arXiv:1609.03769v1 [stat.ML] 13 Sep 2016

Analysis of Kelner and Levin graph sparsification algorithm for a streaming setting

Daniele Calandriello **Alessandro Lazaric** **Michal Valko**
SequeL team, INRIA Lille - Nord Europe, France
{daniele.calandriello, alessandro.lazaric, michal.valko}@inria.fr

Abstract

We derive a new proof to show that the incremental resparsification algorithm proposed by Kelner and Levin [2] produces a spectral sparsifier in high probability. We rigorously take into account the dependencies across subsequent resparsifications using martingale inequalities, fixing a flaw in the original analysis.

1 Introduction

Kelner and Levin [2] introduced a simple single-pass approach to generate a spectral sparsifier of a graph in the semi-streaming setting, where edges are received one at a time. They store only an intermediate, approximate sparsifier and every time a new edge arrives, it is added to it. Whenever the sparsifier gets too large, they apply a resparsification algorithm to reduce its size, without compromising its spectral guarantees.

Although the algorithm is intuitive and simple to implement, the original proof presented in their paper is incomplete, as originally pointed out in Cohen et al. [1]. In particular, Kelner and Levin [2] relies on a concentration inequality for independent random variables, while in the sparsification algorithm the probability of keeping edge e in the sparsifer at step s does depend on whether other edges e' have been included in the sparsifier at previous iterations. This structure introduces subtle statistical dependencies through different iterations of the algorithm, and a more careful analysis is necessary.

In addition to pointing out the problems with the original proof in [2], Cohen et al. [1] introduces a new algorithm to construct a sparsifer in a semi-streaming setting but, differently from the original algorithm in [2], interactions between iterations are avoided because the algorithm proposed in [1] never drops an edge once it is introduced in the sparsifier. Another alternative algorithm, this time capable of dropping included edges, is presented in Pachocki [4] together with a rigorous proof that takes into account all the dependencies between edges and between iterations. While the final result in [4] guarantees that a valid spectral sparisfier is generated at each iteration, the proposed algorithm is still different from the one originally proposed by Kelner and Levin [2].

In this note, we derive an alternative proof for the original Kelner and Levin [2] algorithm, using arguments similar to [4]. In particular, it is possible to formalize and analyze the edge selection process as a martingale, obtaining strong concentration guarantees while rigorously taking into account the dependencies across the iterations of the algorithm.

2 Background

2.1 Notation

We use lowercase letters a for scalars, bold lowercase letters $\mathbf{a}$ for vectors and uppercase bold letters $\mathbf{A}$ for matrices. We write $\mathbf{A} \preceq \mathbf{B}$ for the Löwner ordering of matrices $\mathbf{A}$ and $\mathbf{B}$ when $\mathbf{B} - \mathbf{A}$ is positive semi-definite (PSD).

We denote with $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ an undirected weighted graph with n vertices $\mathcal{V}$ and m edges $\mathcal{E}$. Associated with each edge $e_{i,j} \in \mathcal{E}$ there is a weight $a_{e_{i,j}}$ (shortened a_e) measuring the "distance" between vertex i and vertex j.[1] Throughout the rest of the paper, we assume that the weights a_e are bounded, in particular we assume $a_{\max} = \max_{e\in\mathcal{E}} a_e$ is smaller than 1, $a_{\min} = \min_{e\in\mathcal{E}} a_e$ is strictly greater than 0, and that $\kappa^2 = a_{\max}/a_{\min} = \mathcal{O}(\text{poly}(n))$, which is always true for unweighted graphs. Given two graphs $\mathcal{G}$ and $\mathcal{G}'$ over the same set of nodes $\mathcal{V}$, we denote by $\mathcal{G} + \mathcal{G}'$ the graph obtained by summing the weights of the edges of $\mathcal{G}'$ and $\mathcal{G}$.

Given the weighted adjacency matrix $\mathbf{A}_\mathcal{G}$ and the degree matrix $\mathbf{D}_\mathcal{G}$, the Laplacian of $\mathcal{G}$ is the PSD matrix defined as $\mathbf{L}_\mathcal{G} = \mathbf{D}_\mathcal{G} - \mathbf{A}_\mathcal{G}$. Furthermore, we assume that $\mathcal{G}$ is connected and thus has only one eigenvalue equal to 0 and $\text{Ker}(\mathbf{L}_\mathcal{G}) = \mathbf{1}$. Let $\mathbf{L}_\mathcal{G}^+$ be the pseudoinverse of $\mathbf{L}_\mathcal{G}$, and $\mathbf{L}_\mathcal{G}^{-1/2} = (\mathbf{L}_\mathcal{G}^+)^{1/2}$. For any node $i = 1, \ldots, n$, we denote with $\chi_i \in \mathbb{R}^n$ the indicator vector so that $\mathbf{b}_e = \chi_i - \chi_j$ is the "edge" vector. If we denote with $\mathbf{B}_\mathcal{G}$ the $m \times n$ signed edge-vertex incidence matrix, then the Laplacian matrix can be written as $\mathbf{L}_\mathcal{G} = \sum_e a_e \mathbf{b}_e \mathbf{b}_e^\mathsf{T} = \mathbf{B}_\mathcal{G}^\mathsf{T} \mathbf{E}_\mathcal{G} \mathbf{B}_\mathcal{G}$, where $\mathbf{E}_\mathcal{G}$ is the $m \times m$ diagonal matrix with $\mathbf{E}_\mathcal{G}(e, e) = a_e$.

We indicate with $\mathbf{P} = \mathbf{L}_\mathcal{G}\mathbf{L}_\mathcal{G}^+$ the matrix of the orthogonal projection on the $n-1$ dimensional space ortoghonal to the all one vector $\mathbf{1}$. Since the Laplacian of any connected graph $\mathcal{G}$ has a null space equal to $\mathbf{1}$, then $\mathbf{P}$ is invariant w.r.t. the specific graph $\mathcal{G}$ on n vertices used to defined it. Alternatively, the projection matrix $\mathbf{P}$ can be obtained as $\mathbf{P} = \mathbf{L}_\mathcal{G}^+\mathbf{L}_\mathcal{G} = \mathbf{L}_\mathcal{G}^{-1/2}\mathbf{L}_\mathcal{G}\mathbf{L}_\mathcal{G}^{-1/2}$. Finally, let $\mathbf{v}_e = \sqrt{a_e}\mathbf{L}_\mathcal{G}^{-1/2}\mathbf{b}_e$, then we have $\mathbf{P} = \sum_e \mathbf{v}_e \mathbf{v}_e^\mathsf{T}$.

2.2 Spectral Sparsification in the Semi-Streaming Setting

A graph $\mathcal{H}$ is a spectral sparsifier of $\mathcal{G}$ if the whole spectrum of the original graph is well approximated by using only a small portion of its edges. More formally,

Definition 1. *A $1 \pm \varepsilon$ spectral sparsifier of $\mathcal{G}$ is a graph $\mathcal{H} \subseteq \mathcal{G}$ such that for all $\mathbf{x} \in \mathbb{R}^n$*

$$(1-\varepsilon)\mathbf{x}^\mathsf{T}\mathbf{L}_\mathcal{G}\mathbf{x} \leq \mathbf{x}^\mathsf{T}\mathbf{L}_\mathcal{H}\mathbf{x} \leq (1+\varepsilon)\mathbf{x}^\mathsf{T}\mathbf{L}_\mathcal{G}\mathbf{x}.$$

Spectral sparsifiers store most of the spectral information of the original graph in a very sparse subgraph. Because of this, they are easy to store in memory and are used to provide fast approximation to many quantities that are expensive to compute on the original large graph.

After showing that every graph admits a sparsifier with $\mathcal{O}(n\log(n)/\varepsilon^2)$ edges, Spielman and Srivastava [5] proposed a sampling algorithm to easily construct one using the effective resistance of the edges of $\mathcal{G}$.

Definition 2. *The effective resistance of an edge e in graph $\mathcal{G}$ is defined as $r_e = \mathbf{b}_e^\mathsf{T}\mathbf{L}_\mathcal{G}^+\mathbf{b}_e$. The total weighted sum of effective resistances in a graph is the same for all graphs, and is equal to $\sum_e a_e r_e = \text{Tr}(\mathbf{E}_\mathcal{G}\mathbf{B}_\mathcal{G}\mathbf{L}_\mathcal{G}^+\mathbf{B}_\mathcal{G}^\mathsf{T}) = \text{Tr}(\mathbf{L}_\mathcal{G}\mathbf{L}_\mathcal{G}^+) = n-1$.*

Intuitively, the effective resistance encodes the importance of an edge in preserving the minimum distance between two nodes. If an edge is the only connection between two parts of the graph, its r_e is large. On the other hand, if there are multiple parallel paths across many edges to connect two nodes, the effective resistance of an edge between the two nodes will be small, similarly to actual resistances in parallel in an electrical network. An important consequence of this definition is that adding edges to a graph can only reduce the effective resistance of other edges, because it can only introduce new alternative (parallel) paths in the graph. To prove this formally, consider a graph $\mathcal{G}$ and a new set of edges Γ. Then we have $\mathbf{L}_\mathcal{G} \preceq \mathbf{L}_{\mathcal{G}+\Gamma}$ and therefore $\mathbf{b}_e^\mathsf{T}\mathbf{L}_\mathcal{G}^+\mathbf{b}_e \geq \mathbf{b}_e^\mathsf{T}\mathbf{L}_{\mathcal{G}+\Gamma}^+\mathbf{b}_e$.

Spielman and Srivastava [5] proved that sampling the edges of $\mathcal{G}$ with replacement using a distribution proportional to their effective resistance produces a spectral sparsifier $\mathcal{H}$ of size $\mathcal{O}(n\log(n)/\varepsilon^2)$ with high probability. The main issue of this approach is that we want to compute a sparsifier to avoid storing the whole Laplacian, but we need to store and (pseudo-)invert the Laplacian to compute exact effective resistances to construct the sparsifier. Spielman and Srivastava [5] showed that this issue can be resolved by computing sufficiently accurate approximation of the effective resistances.

[1] The graph $\mathcal{G}$ can be either constructed from raw data (e.g., building a k-nn graph with an exponential kernel) or it can be provided directly as input (e.g., in social networks).

Algorithm 1 Kelner and Levin [2] stream sparsification algorithm.

Require: Graph $\mathcal{G}$, weights a_e for all edges in $\mathcal{G}$.
Ensure: $\mathcal{H}_\tau$, a $1 \pm \varepsilon$ sparsifier of $\mathcal{G}$

1: Set space budget $N = 40\alpha^2 n \log^2(3\kappa m/\delta)/\varepsilon^2$
2: Partition $\mathcal{G}$ into $\tau = \lceil m/N \rceil$ blocks $\{\Gamma_s\}_{s=0}^{\tau}$ such that $\mathcal{G} = \mathcal{G}_\tau = \sum_{s=0}^{\tau} \Gamma_s$
3: Initialize $\mathcal{H}_0 = \emptyset$, $\{\widetilde{p}_{e,0}\} = \emptyset$
4: **for** all $s \in [1, \ldots, \tau]$ **do**
5: Use Algorithm 2 to compute an updated sparsifier $\mathcal{H}_s$ of $\mathcal{G}_s = \mathcal{G}_{s-1} + \Gamma_s$ and new probabilities $\{\widetilde{p}_{s,e}\}$ for all edges in $\mathcal{H}_s$.
6: **end for**
7: Return $\mathcal{H}_\tau$

Algorithm 2 Kelner and Levin [2] resparsification algorithm at step s.

Require: $\mathcal{H}_{s-1}$,Γ_s, the set of previous probabilities $\{\widetilde{p}_{s-1,e}\}$ for all edges in $\mathcal{H}_{s-1}$ and the weights a_e for all edges in $\mathcal{H}_{s-1} + \Gamma_s$
Ensure: $\mathcal{H}_s$, a $1 \pm \varepsilon$ sparsifier of $\mathcal{G}_s = \mathcal{G}_{s-1} + \Gamma_s$ and new probabilities $\widetilde{p}_{s,e}$.

1: Compute estimates of all $\widetilde{r}_{s,e}$ in $\mathcal{H}_{s-1} + \Gamma_s$ using a fast SDD solver ([2, Theorem 3])
2: Compute new probabilities $\widetilde{p}_{s,e} = (a_e \widetilde{r}_{s,e})/(\alpha(n-1))$
3: **for** all edges $e \in \mathcal{H}_{s-1}$ **do**
4: $\widetilde{p}_{s,e} \leftarrow \min\{\widetilde{p}_{s-1,e}, \widetilde{p}_{s,e}\}$
5: **end for**
6: Initialize $\mathcal{H}_s = \emptyset$
7: **for** all edges $e \in \mathcal{H}_{s-1}$ **do**
8: With probability $\widetilde{p}_{s,e}/\widetilde{p}_{s-1,e}$ add edge e to $\mathcal{H}_s$ with weight $a_e/(N\widetilde{p}_{s,e})$
9: **end for**
10: **for** all edges $e \in \Gamma_s$ **do**
11: **for** $i = 1$ to N **do**
12: With probability $\widetilde{p}_{s,e}$ add a copy of edge e to $\mathcal{H}_s$ with weight $a_e/(N\widetilde{p}_{s,e})$
13: **end for**
14: **end for**

Definition 3. *An approximate effective resistance $\widetilde{r}_e$ is called α-accurate for $\alpha \geq 1$ if it satisfies*

$$\frac{1}{\alpha} r_e \leq \widetilde{r}_e \leq \alpha r_e.$$

In particular, Spielman and Srivastava [5, Corollary 6] showed that batch sampling $\mathcal{O}(\alpha^2 n \log(n)/\varepsilon^2)$ edges proportionally to their α-accurate approximate effective resistances is enough to guarantee that the resulting graph is a $1 \pm \varepsilon$-sparsifier. Building on this result, Kelner and Levin [2] propose a sequential algorithm (summarized in Alg. 1) that can emulate the batch sampling of [5] in a semi-streaming setting and incrementally construct a sparsifier, without having to fully store and invert the input Laplacian.

In a semi-streaming setting the graph $\mathcal{G}$ is split int $\tau = \lceil m/N \rceil$ blocks Γ_s, with $s \in [1, 2, \ldots, \tau]$, where each block is a subset of N edges such that $\mathcal{G} = \mathcal{G}_\tau = \mathcal{G}_{\tau-1} + \Gamma_\tau = \Gamma_1 + \Gamma_2 + \cdots + \Gamma_\tau$. Associated with each of the $\mathcal{G}_s$ partial graph, we can define its respective effective resistances $r_{s,e}$, and the sampling probabilities $p_{s,e} = (a_e r_{s,e})/(n-1)$. Starting ($s = 0$) from an empty sparsifier $\mathcal{H}_0 = \emptyset$, the algorithm alternates between two phases. In the first phase, the algorithm reads N edges from the stream to build a new block Γ_s, and it combines it with the previous sparsifier $\mathcal{H}_{s-1}$ to construct $\mathcal{H}_{s-1} + \Gamma_s$. This phase sees an increase in memory usage, because the algorithm needs to store the newly arrived edges in addition to the sparsifier. In the second phase the graph $\mathcal{H}_{s-1} + \Gamma_s$ is used together with a fast SDD solver to compute α-accurate estimates $\widetilde{r}_{s,e}$ of the effective resistance of all the edges in it. These approximate effective resistances $\widetilde{r}_{s,e}$ are used to compute approximate probabilities $\widetilde{p}_{s,e}$, and according to these approximate probabilities each of the edges in $\mathcal{H}_{s-1}$ and Γ_s is added to the new sparsifier $\mathcal{H}_s$ or discarded forever freeing up memory. Choosing carefully the size of the blocks Γ_s to be close to the size of the sparsifiers $\mathcal{H}_s$ allows the algorithm to run efficiently in

Algorithm 3 Equivalent formulation of Kelner and Levin's algorithm.

Require: Graph $\mathcal{G}$, weights a_e for all edges in $\mathcal{G}$.
Ensure: $\mathcal{H}_m$, a $1 \pm \varepsilon$ sparsifier of $\mathcal{G}$

1: Set space budget $N = 40\alpha^2 n \log^2(3\kappa m/\delta)/\varepsilon^2$
2: Set $\widehat{z}_{s,e,j} = 1, \widetilde{p}_{s,e} = 1$ for all $e = 1, \dots, m$, $s = 0, \dots, m$, $j = 1, \dots, N$
3: **for** $s = 1, \dots, m$ **do** ▷ Stream over all edges in $\mathcal{G}$
4: **for** $e = 1, \dots, m$ **do**
5: **if** $e \leq s$ **then**
6: Compute $\widetilde{r}_{s,e}$ using $\mathcal{H}_{s-1}$ and fast SDD solver
7: Compute $\widetilde{p}_{s,e} = \min\left\{\frac{a_e \widetilde{r}_{s,e}}{\alpha(n-1)}, \widetilde{p}_{s-1,e}\right\}$
8: **end if**
9: **for** $j = 1, \dots, N$ **do** ▷ Loop over N trials (i.e., copies)
10: **if** $\widehat{z}_{s-1,e,j} = 1$ **then**
11: Sample $\widehat{z}_{s,e,j} \sim \mathcal{B}\left(\frac{\widetilde{p}_{s,e}}{\widetilde{p}_{s-1,e}}\right)$
12: **else**
13: $\widehat{z}_{s,e,j} = 0$
14: **end if**
15: **end for**
16: **end for**
17: Set $\mathcal{H}_s = \emptyset$
18: **for** $e = 1, \dots, s$, $j = 1, \dots, N$ **do** ▷ Construction of the new sparsifier
19: **if** $\widehat{z}_{s,e,j} = 1$ **then**
20: Add edge e to $\mathcal{H}_s$ with weight $a_e/(N\widetilde{p}_{s,e})$
21: **end if**
22: **end for**
23: **end for**

a small fixed space and produce a valid sparsifier at the end of each iteration.[2] For more details on the time complexity analysis and the implementation details on how to obtain α-approximate effective resistances using a valid sparsifier and a fast SDD solver we refer to the original paper [2].

The main result of Kelner and Levin [2] is the following theorem.

Theorem 1. *Let $\mathcal{H}_s$ be the sparsifier returned by Algorithm 1 after s resparsifications (after streaming the first s blocks). If $\alpha \leq \sqrt{\kappa n/3}$ and $\kappa^2 = \mathcal{O}(poly(n))$, with probability $1 - \delta$ all sparsifiers $\mathcal{H}_s$ from the beginning of the algorithm to its end ($\forall s \in [1, \dots, \tau]$) are valid $(1 \pm \varepsilon)$ spectral sparsifiers of their corresponding partial graph $\mathcal{G}_s$ and the number of edges in each of the $\mathcal{H}_s$ sparsifiers is $\mathcal{O}(n \log^2(\kappa m)/\varepsilon^2) = \mathcal{O}(n \log^2(n)/\varepsilon^2)$.*

At the core of the original proof of this theorem, Kelner and Levin rely on a concentration inequality for independent random variables in [8]. Unfortunately, it is not possible to directly use this result since the probability that an edge e is included in $\mathcal{H}_s$ does indeed depend on all the edges that were included in $\mathcal{H}_{s-1}$ through the computation of $\widetilde{r}_{s,e}$. The sparsifier $\mathcal{H}_{s-1}$ in turn is generated from $\mathcal{H}_{s-2}$ and so on. As a result, the probability that an edge e is present in the final graph $\mathcal{H}_\tau$ is strictly dependent on the other edges. In the following we rigorously take into account the interactions across iterations and provide a new proof for Theorem 1 which confirms its original statement, thus proving the correctness of [2]'s algorithm.

3 Proof

Step 1 (the *theoretical* algorithm). In order to simplify the analysis, we introduce an equivalent formulation of Alg. 1. In Alg. 3 we consider the case where the blocks Γ_s contain only a single edge and the algorithm performs m resparsifications over the course of the whole stream of m edges (loop at line 3). This is a wasteful approach and more practical methods (such as Alg. 1) choose to

[2] Throughout this note, we consider that the decomposition of $\mathcal{G}$ into blocks is such that all intermediate graphs $\mathcal{G}_s$ are fully connected. Whenever this is not the case, the algorithm should be adjusted to run separately on all the components of the graph.

resparsify only when Γ_s is larger than $\mathcal{H}_{s-1}$ in order to save time without increasing the asymptotic space complexity of the algorithm. Nonetheless, the single edge setting can be used to emulate any larger choice for the block size Γ_s, and therefore we only need to prove our result in this setting for it to hold in any other case.

In Alg. 3, we denote by $\widehat{z}_{s,e,j}$ the Bernoulli random variable that indicates whether copy j of edge e is present in the sparsifier $\mathcal{H}_s$ at step s. While an edge e can be present in $\mathcal{H}_s$ only if $e \leq s$, we initialize $\widehat{z}_{s,e,j} = 1$ for all $e > s$ for notational convenience. The way these variables are generated is equivalent to lines 7-14 in Alg. 2 so that at iteration s, each copy of an edge e already in the sparsifier $\mathcal{H}_{s-1}$ is kept with probability $\widetilde{p}_{s,e}/\widetilde{p}_{s-1,e}$, while a copy of the new edge s is added with probability $\widetilde{p}_{s,e}$. For any edge $e > s$ (i.e., not processed yet in Alg. 2) we initialize $\widetilde{p}_{s,e} = 1$ and thus the sampling in line 10 of Alg. 3 always returns $\widehat{z}_{s,e,j} = 1$. Since edges are added with weights $a_e/(N\widetilde{p}_{s,e})$, after processing $s \leq m$ edges, the Laplacian of the sparsifier $\mathcal{H}_s$ can be written as

$$\mathbf{L}_{\mathcal{H}_s} = \sum_{e=1}^{s} \sum_{j=1}^{N} \frac{a_e}{N\widetilde{p}_{s,e}} \widehat{z}_{s,e,j} \mathbf{b}_e \mathbf{b}_e^{\mathsf{T}}.$$

Step 2 (filtration). A convenient way to treat the indicator variables $\widehat{z}_{s,e,j}$ is to define them recursively as

$$\widehat{z}_{s,e,j} \stackrel{\text{def}}{=} \mathbb{I}\left\{ u_{s,e,j} \leq \frac{\widetilde{p}_{s,e}}{\widetilde{p}_{s-1,e}} \right\} \widehat{z}_{s-1,e,j},$$

where $u_{s,e,j} \sim \mathcal{U}(0,1)$ is a uniform random variable used to compute the $\widetilde{p}_{s,e}/\widetilde{p}_{s-1,e}$ coin flip, and $\widetilde{p}_{s,e}$ are the approximate probabilities computed at step s according to the definition in Algorithm 2, using the SDD solver, the sparsifier $\mathcal{H}_{s-1}$ and the new edges Γ_s. This formulation allows us to define the stochastic properties of variables $\widehat{z}_{s,e,j}$ in a convenient way. We first arrange the indices s, e, and j into a linear index $r = \{s, e, j\}$ in the range $[1, \ldots, m^2N]$, obtained as $r = \{s, e, j\} = (s-1)mN + (e-1)N + j$. Following the structure of Alg. 3, the linearization wraps first when j hits its limit, and then when e and finally s do the same, such that for any $s < m$, $e < m$, and $j < N$, we have

$$\begin{aligned}
\{s, e, j\} + 1 &= \{s, e, j+1\}, \\
\{s, e, N\} + 1 &= \{s, e+1, 1\}, \\
\{s, m, N\} + 1 &= \{s+1, 1, 1\}.
\end{aligned}$$

It is easy to see that the checkpoints $\{s, m, N\}$ correspond to a full iteration (Alg. 3, line 3) of the algorithm. Let $\mathcal{F}_{\{s,e,j\}}$ be the filtration containing all the realizations of the uniform random variables $u_{s,e,j}$ up to the step $\{s, e, j\}$, that is $\mathcal{F}_{\{s,e,j\}} = \{u_{s',e',j'}, \forall \{s', e', j'\} \leq \{s, e, j\}\}$. Again, we notice that $\mathcal{F}_{\{s,m,N\}}$ defines the state of the algorithm after completing iteration s. Since $\widetilde{r}_{s,e}$ and $\widetilde{p}_{s,e}$ are computed at the beginning of iteration s using the sparisfier $\mathcal{H}_{s-1}$, they are fully determined by $\mathcal{F}_{\{s-1,m,N\}}$. Furthermore, since $\mathcal{F}_{\{s-1,m,N\}}$ also defines the values of all indicator variables $\widehat{z}_{s',e,j}$ up to $\widehat{z}_{s-1,e,j}$ for any e and j, we have that all the Bernoulli variables $\widehat{z}_{s,e,j}$ at iteration s are conditionally independent given $\mathcal{F}_{\{s-1,m,N\}}$. In other words, we have that for any e', and j' such that $\{s, 1, 1\} \leq \{s, e', j'\} < \{s, e, j\}$ the following random variables are equal in distribution

$$\widehat{z}_{s,e,j} \big| \mathcal{F}_{\{s,e',j'\}} = \widehat{z}_{s,e,j} \big| \mathcal{F}_{\{s-1,m,N\}} \sim \mathcal{B}\Big(\frac{\widetilde{p}_{s,e}}{\widetilde{p}_{s-1,e}}\Big) \tag{1}$$

and for any e', and j' such that $\{s, 1, 1\} \leq \{s, e', j'\} \leq \{s, m, N\}$ and $\{s, e, j\} \neq \{s, e', j'\}$

$$\widehat{z}_{s,e,j} \big| \mathcal{F}_{\{s-1,m,N\}} \perp \widehat{z}_{s,e',j'} \big| \mathcal{F}_{\{s-1,m,N\}}. \tag{2}$$

Step 3 (the projection error). While our objective is to show that $\mathcal{H}_m$ is a $1 \pm \varepsilon$-sparsifier, following [2] we study the related objective of defining an approximate projection matrix that is close in ℓ_2-norm to the original projection matrix $\mathbf{P}$. In fact, the two objectives are strictly related as shown in the following proposition.

Proposition 1 ([2]). *Given $\mathbf{v}_e = \sqrt{a_e}\mathbf{L}_{\mathcal{G}}^{-1/2}\mathbf{b}_e$, let $\mathcal{I} = [e_1, e_2, \ldots, e_N] \in \mathcal{E}^N$ be a subset of edges of $\mathcal{G}$ and $\widetilde{\mathbf{P}} = \sum_{i=1}^{N} w_i \mathbf{v}_{e_i} \mathbf{v}_{e_i}^{\mathsf{T}}$ an approximate projection matrix with weights $\{w_i\}_{i=1}^{N}$. If the weights*

are such that

$$\left\|\mathbf{P}-\sum_{i=1}^{N} w_i \mathbf{v}_{e_i}\mathbf{v}_{e_i}^{\mathsf{T}}\right\|_2 = \left\|\mathbf{P}-\widetilde{\mathbf{P}}\right\|_2 \leq \varepsilon, \tag{3}$$

then the graph $\mathcal{H}$ obtained by adding edges $e_i \in \mathcal{I}$ with weights $a_{e_i}w_i$ is a $(1\pm\varepsilon)$-sparsifier of $\mathcal{G}$.

Using the notation of Alg. 3, the approximate projection matrix is defined as

$$\widetilde{\mathbf{P}} \stackrel{\text{def}}{=\!=} \frac{1}{N}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{\widehat{z}_{m,e,j}}{\widetilde{p}_{m,e}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}},$$

and thus the previous proposition suggests that to prove that the graph $\mathcal{H}_m$ returned by Algorithm 1 after m steps is a sparsifier, it is sufficient to show that

$$\left\|\mathbf{P}-\widetilde{\mathbf{P}}\right\|_2 = \left\|\frac{1}{N}\sum_{j=1}^{N}\sum_{e=1}^{m}\left(1-\frac{\widehat{z}_{m,e,j}}{\widetilde{p}_{m,e}}\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right\|_2 \leq \varepsilon.$$

In order to study this quantity, we need to analyze how the projection error evolves over iterations. To this end, we introduce term $\widehat{\mathbf{Y}}_{\{s,e,j\}}$ which denotes the projection error at the end of step $\{s,e,j\}$ of Algorithm 3.

$$\begin{aligned}\widehat{\mathbf{Y}}_{\{s,e,j\}} \stackrel{\text{def}}{=\!=} &\underbrace{\frac{1}{N}\sum_{k=1}^{e-1}\sum_{l=1}^{N}\left(1-\frac{\widehat{z}_{s,k,l}}{\widetilde{p}_{s,k}}\right)\mathbf{v}_k\mathbf{v}_k^{\mathsf{T}}}_{\text{edges already processed at step } s}\\ &+\underbrace{\frac{1}{N}\left(\sum_{l=1}^{j}\left(1-\frac{\widehat{z}_{s,e,l}}{\widetilde{p}_{s,e}}\right)+\sum_{l=j+1}^{N}\left(1-\frac{\widehat{z}_{s-1,e,l}}{\widetilde{p}_{s-1,e}}\right)\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}}_{\text{copy } j \text{ of edge } e \text{ getting processed at step } s \text{ (Alg. 3, line 11)}}\\ &+\underbrace{\frac{1}{N}\sum_{k=e+1}^{m}\sum_{l=1}^{N}\left(1-\frac{\widehat{z}_{s-1,k,l}}{\widetilde{p}_{s-1,k}}\right)\mathbf{v}_k\mathbf{v}_k^{\mathsf{T}}}_{\text{edges still not processed at step } s}\end{aligned}$$

Notice that setting $\widehat{z}_{s,e,j}=1$ and $\widetilde{p}_{s,e}=1$ for any $e>s$ implies that the edges that have not been processed yet do not contribute to the projection error. Finally, notice that at the end of the algorithm we have $\widehat{\mathbf{Y}}_{\{m,m,N\}} = \mathbf{P}-\widetilde{\mathbf{P}}$, which quantifies the error of the output of Algorithm 1.

We are now ready to restate Theorem 1 in a more convenient way as

$$\mathbb{P}\left(\exists s\in\{1,\dots,m\}: \underbrace{\|\widehat{\mathbf{Y}}_{\{s,m,N\}}\|\geq\varepsilon}_{A_s} \cup \underbrace{\sum_{e=1}^{s}\sum_{j=1}^{N}\widehat{z}_{s,e,j}\geq 3N}_{B_s}\right)\leq\delta,$$

where the first event A_s refers to the case when for any $s\in[1,\dots,m]$ the intermediate graph $\mathcal{H}_s$ fails to be a valid sparsifier and the second event B_s considers the event when the memory requirement is not met (i.e., too many edges are kept in the sparsifier $\mathcal{H}_s$).

To prove the statement, we decompose the probability of failure as follows.

$$\mathbb{P}\Big(\exists s \in \{1,\dots,m\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \geq \varepsilon \cup \sum_{e=1}^{s}\sum_{j=1}^{N} \widehat{z}_{s,e,j} \geq 3N\Big) = \mathbb{P}\Big(\bigcup_{s=1}^{m} A_s \cup B_s\Big)$$

$$= \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} A_s\right\} \cup \left\{\bigcup_{s=1}^{m} B_s\right\}\right)$$

$$= \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} A_s\right\}\right) + \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} B_s\right\}\right) - \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} A_s\right\} \cap \left\{\bigcup_{s=1}^{m} B_s\right\}\right)$$

$$= \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} A_s\right\}\right) + \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} B_s\right\} \cap \left\{\bigcup_{s=1}^{m} A_s\right\}^{\mathsf{c}}\right)$$

$$= \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} A_s\right\}\right) + \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} B_s\right\} \cap \left\{\bigcap_{s=1}^{m} A_s^{\mathsf{c}}\right\}\right)$$

$$= \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} A_s\right\}\right) + \mathbb{P}\left(\bigcup_{s=1}^{m}\left\{B_s \cap \left\{\bigcap_{s'=1}^{m} A_{s'}^{\mathsf{c}}\right\}\right\}\right)$$

Taking the last formulation and replacing the definitions of A_s and B_s, we get

$$\mathbb{P}\Big(\exists s \in \{1,\dots,m\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \geq \varepsilon \cup \sum_{e=1}^{s}\sum_{j=1}^{N} \widehat{z}_{s,e,j} \geq 3N\Big)$$

$$= \mathbb{P}\left(\left\{\bigcup_{s=1}^{m} \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \geq \varepsilon\right\}\right) + \mathbb{P}\left(\bigcup_{s=1}^{m}\left\{\sum_{e=1}^{s}\sum_{j=1}^{N} \widehat{z}_{s,e,j} \geq 3N \cap \left\{\bigcap_{s'=1}^{m} \|\widehat{\mathbf{Y}}_{\{s',m,N\}}\| \leq \varepsilon\right\}\right\}\right)$$

$$= \mathbb{P}\Big(\exists s \in \{1,\dots,m\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \geq \varepsilon\Big)$$

$$+ \mathbb{P}\left(\exists s \in \{1,\dots,m\} : \sum_{e=1}^{s}\sum_{j=1}^{N} \widehat{z}_{s,e,j} \geq 3N \cap \Big\{\forall s' \in \{1,\dots,m\} : \|\widehat{\mathbf{Y}}_{\{s',m,N\}}\| \leq \varepsilon\Big\}\right)$$

$$\leq \sum_{s=1}^{m} \mathbb{P}\Big(\|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \geq \varepsilon\Big)$$

$$+ \sum_{s=1}^{m} \mathbb{P}\left(\sum_{e=1}^{s}\sum_{j=1}^{N} \widehat{z}_{s,e,j} \geq 3N \cap \Big\{\forall s' \in \{1,\dots,m\} : \|\widehat{\mathbf{Y}}_{\{s',m,N\}}\| \leq \varepsilon\Big\}\right)$$

$$\leq \sum_{s=1}^{m} \mathbb{P}\Big(\|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \geq \varepsilon\Big)$$

$$+ \sum_{s=1}^{m} \mathbb{P}\left(\sum_{e=1}^{s}\sum_{j=1}^{N} \widehat{z}_{s,e,j} \geq 3N \cap \Big\{\forall s' \in \{1,\dots,s\} : \|\widehat{\mathbf{Y}}_{\{s',m,N\}}\| \leq \varepsilon\Big\}\right) \tag{4}$$

Step 4 (putting everything together). In the following sections, we prove the two main lemmas of this note, where we bound the probability of returning a non-spectral sparsifier and the probability of exceeding too much the budget limit N. In particular, we derive the two following results.

Lemma 1.

$$\mathbb{P}\Big(\|\widehat{\mathbf{Y}}_{\{t,m,N\}}\| \geq \varepsilon\Big) \leq \frac{\delta}{2m}$$

Lemma 2.

$$\mathbb{P}\left(\sum_{e=1}^{t}\sum_{j=1}^{N} \widehat{z}_{t,e,j} \geq 3N \cap \Big\{\forall\, s \in \{1,\dots,t\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \leq \varepsilon\Big\}\right) \leq \frac{\delta}{2m}$$

Combining the two lemmas into Eq. 4, we prove Thm .1 for an algorithm that resparsifies every time a new edge arrives ($|\Gamma_s| = 1$). Extending the proof to the case when multiple edges are stored in Γ_s before a new resparsification happens is straightforward. In the proofs of Lemma 1 and 2 the fact that an edge is unseen (not streamed yet) is represented by deterministically setting its $\widetilde{p}_{s,e}$ to 1, while the estimates for seen edges are computed based on the graph. To represent the arrival of an edge e at time t we simply start updating its $\widetilde{p}_{t,e}$. To take into account τ resparsifications of large blocks instead of m resparsifications of single-edge blocks it is sufficient to start updating multiple $\widetilde{p}_{t,e}$ at the same step. The rest of the analysis remains unchanged.

4 Proof of Lemma 1 (bounding $\widehat{\mathbf{Y}}_{\{s,m,N\}}$)

Step 1 (freezing process). We first restate a proposition on the accuracy of the effective resistance estimates.

Proposition 2. *At iteration s the approximated effective resistance $\widetilde{r}_{s,e}$ of an edge e in $\mathcal{H}_{s-1} + \Gamma_s$ is computed using $\mathcal{H}_{s-1} + \Gamma_s$ and the SDD solver. If $\mathcal{H}_{s-1}$ is a valid $1 \pm \varepsilon$-sparsifier of $\mathcal{G}_{s-1}$, then $\widetilde{r}_{s,e}$ is α-accurate.*

Given α-accurate effective resistances, the approximate probabilities $\widetilde{p}_{s,e}$ are defined as

$$\widetilde{p}_{s,e} = \min\left\{\frac{a_e \widetilde{r}_{s,e}}{\alpha(n-1)}, \widetilde{p}_{s-1,e}\right\}.$$

As pointed out in Proposition 2, the main issue is that whenever $\mathcal{H}_{s-1}$ is not a valid sparsifier of $\mathcal{G}_{s-1}$, the approximate probabilities $\widetilde{p}_{s,e}$ returned by the fast SDD solver are not guaranteed to be α-accurate approximations of the true probabilities $p_{s,e}$. While the overall algorithm may fail in generating a valid sparsifier at some intermediate iteration and yet return a valid sparsifier at the end, we consider an alternative (more pessimistic) process which is "frozen" as soon as it constructs an invalid sparsifier. Consider an alternative process $\mathbf{Y}_{\{s,e,j\}}$ based on the following definition of approximate probabilities

$$\overline{p}_{s,e} = \widetilde{p}_{s,e}\mathbb{I}\left\{\|\mathbf{Y}_{\{s-1,m,N\}}\| \leq \varepsilon\right\} + \overline{p}_{s-1,e}\mathbb{I}\left\{\|\mathbf{Y}_{\{s-1,m,N\}}\| \geq \varepsilon\right\},$$

where by Proposition 1, the condition $\|\mathbf{Y}_{\{s-1,m,N\}}\| \leq \varepsilon$ is equivalent to requiring that $\mathcal{H}_{s-1}$ is a valid sparsifier. This new formulation represents a variant of our algorithm that can detect if the previous iteration failed to construct a graph that is guaranteed to be a sparsifier. When this failure happens, the whole process is frozen and continues until the end without updating anything. Then we redefine the indicator variable $z_{s,e,j}$ dependent on $\overline{p}_{s,e}$ as

$$z_{s,e,j} = \mathbb{I}\left\{u_{s,e,j} \leq \frac{\overline{p}_{s,e}}{\overline{p}_{s-1,e}}\right\} z_{s-1,e,j},$$

and then the projection error process based on them becomes

$$\begin{aligned}
\mathbf{Y}_{\{s,e,j\}} = &\frac{1}{N}\sum_{k=1}^{e-1}\sum_{l=1}^{N}\left(1 - \frac{z_{s,k,l}}{\overline{p}_{s,k}}\right)\mathbf{v}_k\mathbf{v}_k^\mathsf{T} \\
&+ \frac{1}{N}\left(\sum_{l=1}^{j}\left(1 - \frac{z_{s,e,l}}{\overline{p}_{s,e}}\right) + \sum_{l=j+1}^{N}\left(1 - \frac{z_{s-1,e,l}}{\overline{p}_{s-1,e}}\right)\right)\mathbf{v}_e\mathbf{v}_e^\mathsf{T} \\
&+ \frac{1}{N}\sum_{k=e+1}^{m}\sum_{l=1}^{N}\left(1 - \frac{z_{s-1,k,l}}{\overline{p}_{s-1,k}}\right)\mathbf{v}_k\mathbf{v}_k^\mathsf{T}.
\end{aligned}$$

We can see that whenever $\|\mathbf{Y}_{\{s,m,N\}}\| \geq \varepsilon$ at step s, for all successive steps s' we have $z_{s',e,j} = z_{s,e,j}$, or in other words we never drop or add a new edge and never change their weights, since $\overline{p}_{s,e}$ is constant. Consequently, if any of the intermediate elements of the sequence violates the condition $\|\mathbf{Y}_{\{s,e,j\}}\| \leq \varepsilon$, the last element will violate it too. For the rest, the sequence behaves exactly like $\widehat{\mathbf{Y}}_{\{s,e,j\}}$. Therefore,

$$\mathbb{P}\left(\|\widehat{\mathbf{Y}}_{\{t,m,N\}}\| \geq \varepsilon\right) \leq \mathbb{P}\Big(\|\mathbf{Y}_{\{t,m,N\}}\| \geq \varepsilon\Big).$$

Step 2 (martingale process). We now proceed by studying the process $\{\mathbf{Y}_{\{s,e,j\}}\}$ and showing that it is a bounded martingale. The sequence difference process $\{\mathbf{X}_{\{s,e,j\}}\}$ is defined as $\mathbf{X}_{\{s,e,j\}} = \mathbf{Y}_{\{s,e,j\}} - \mathbf{Y}_{\{s,e,j\}-1}$, that is

$$\mathbf{X}_{\{s,e,j\}} = \frac{1}{N}\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s,e,j}}{\overline{p}_{s,e}}\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}.$$

In order to show that $\mathbf{Y}_{\{s,e,j\}}$ is a martingale, it is sufficient to verify the following (equivalent) conditions

$$\mathbb{E}\left[\mathbf{Y}_{\{s,e,j\}} \;\middle|\; \mathcal{F}_{\{s,e,j\}-1}\right] = \mathbf{Y}_{\{s,e,j\}-1} \;\Leftrightarrow\; \mathbb{E}\left[\mathbf{X}_{\{s,e,j\}} \;\middle|\; \mathcal{F}_{\{s,e,j\}-1}\right] = \mathbf{0}.$$

We begin by inspecting the conditional random variable $\mathbf{X}_{\{s,e,j\}}|\mathcal{F}_{\{s,e,j\}-1}$. Given the definition of $\mathbf{X}_{\{s,e,j\}}$, the conditioning on $\mathcal{F}_{\{s,e,j\}-1}$ determines the values of $z_{s-1,e,j}$ and the approximate probabilities $\overline{p}_{s-1,e}$ and $\overline{p}_{s,e}$. In fact, remember that these quantities are fully determined by the realizations in $\mathcal{F}_{\{s-1,m,N\}}$ which are contained in $\mathcal{F}_{\{s,e,j\}-1}$. As a result, the only stochastic quantity in $\mathbf{X}_{\{s,e,j\}}$ is the variable $z_{s,e,j}$. Specifically, if $\|\mathbf{Y}_{\{s-1,m,N\}}\| \geq \varepsilon$, then we have $\overline{p}_{s,e} = \overline{p}_{s-1,e}$ and $z_{s,e,j} = z_{s-1,e,j}$ (the process is stopped), and the martingale requirement $\mathbb{E}\left[\mathbf{X}_{\{s,e,j\}} \;\middle|\; \mathcal{F}_{\{s,e,j\}-1}\right] = \mathbf{0}$ is trivially satisfied. On the other hand, if $\|\mathbf{Y}_{\{s-1,m,N\}}\| \leq \varepsilon$ we have,

$$\begin{aligned}
&\mathop{\mathbb{E}}_{u_{s,e,j}}\left[\frac{1}{N}\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s,e,j}}{\overline{p}_{s,e}}\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}} \;\middle|\; \mathcal{F}_{\{s,e,j\}-1}\right] \\
&\quad= \frac{1}{N}\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s-1,e,j}}{\overline{p}_{s,e}}\,\mathbb{E}\left[\mathbb{I}\left\{u_{s,e,j} \leq \frac{\overline{p}_{s,e}}{\overline{p}_{s-1,e}}\right\} \;\middle|\; \mathcal{F}_{\{s,e,j\}-1}\right]\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}} \\
&\quad= \frac{1}{N}\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s-1,e,j}}{\overline{p}_{s,e}}\frac{\overline{p}_{s,e}}{\overline{p}_{s-1,e}}\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}} = \mathbf{0},
\end{aligned}$$

where we use the recursive definition of $z_{s,e,j}$ and the fact that $u_{s,e,j}$ is a uniform random variable in $[0,1]$. This proves that $\mathbf{Y}_{\{s,e,j\}}$ is indeed a martingale. We now compute an upper-bound R on the norm of the values of the difference process as

$$\begin{aligned}
&\|\mathbf{X}_{\{s,e,j\}}\| = \frac{1}{N}\left|\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s,e,j}}{\overline{p}_{s,e}}\right)\right|\|\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\| = \frac{1}{N}\left|\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s,e,j}}{\overline{p}_{s,e}}\right)\right|\|\mathbf{v}_e\|^2 \\
&\leq \frac{1}{N}\frac{a_e r_{m,e}}{\overline{p}_{s,e}} \leq \frac{1}{N}\frac{\alpha^2 a_e r_{m,e}}{p_{s,e}} \leq \frac{1}{N}\frac{\alpha^2 a_e r_{m,e}}{p_{m,e}} = \frac{1}{N}\frac{\alpha^2(n-1)a_e r_{m,e}}{a_e r_{m,e}} = \frac{\alpha^2(n-1)}{N} \stackrel{\text{def}}{=\!=} R,
\end{aligned}$$

where we use the fact that if, at step s, $\|\mathbf{Y}_{\{s-1,m,N\}}\| \leq \varepsilon$, then the approximate $\widetilde{r}_{s,e}$ are α-accurate by Proposition 2 and thus by definition of $p_{s,e}$,

$$\frac{p_{s,e}}{\alpha^2} \leq \overline{p}_{s,e} \leq p_{s,e}.$$

If instead, $\|\mathbf{Y}_{\{s-1,m,N\}}\| \geq \varepsilon$, the process is stopped and $\|\mathbf{X}_{\{s,e,j\}}\| = \|\mathbf{0}\| = 0 \leq R$.

Step 3 (martingale concentration inequality). We are now ready to use a Freedman matrix inequality from [6] to bound the norm of $\mathbf{Y}$.

Proposition 3 (Theorem 1.2 [6])**.** *Consider a matrix martingale $\{\mathbf{Y}_k : k = 0, 1, 2, \dots\}$ whose values are self-adjoint matrices with dimension d, and let $\{\mathbf{X}_k : k = 1, 2, 3, \dots\}$ be the difference sequence. Assume that the difference sequence is uniformly bounded in the sense that*

$$\|\mathbf{X}_k\|_2 \leq R \quad \textit{almost surely} \quad \textit{for } k = 1, 2, 3, \dots.$$

Define the predictable quadratic variation process of the martingale as

$$\mathbf{W}_k \stackrel{\text{def}}{=\!=} \sum_{j=1}^{k} \mathbb{E}\left[\mathbf{X}_j^2 \;\middle|\; \{\mathbf{X}_s\}_{s=0}^{j-1}\right], \quad \textit{for } k = 1, 2, 3, \dots.$$

Then, for all $\varepsilon \geq 0$ and $\sigma^2 > 0$,

$$\mathbb{P}\left(\exists k \geq 0 : \|\mathbf{Y}_k\|_2 \geq \varepsilon \;\cap\; \|\mathbf{W}_k\| \leq \sigma^2\right) \leq 2d \cdot \exp\left\{-\frac{\varepsilon^2/2}{\sigma^2 + R\varepsilon/3}\right\}.$$

In order to use the previous inequality, we develop the probability of error for any fixed t as

$$\begin{aligned}
&\mathbb{P}\left(\|\widehat{\mathbf{Y}}_{\{t,m,N\}}\| \geq \varepsilon\right) \\
&\leq \mathbb{P}\left(\|\mathbf{Y}_{\{t,m,N\}}\| \geq \varepsilon\right) \\
&= \mathbb{P}\left(\|\mathbf{Y}_{\{t,m,N\}}\| \geq \varepsilon \cap \|\mathbf{W}_{\{t,m,N\}}\| \leq \sigma^2\right) + \mathbb{P}\left(\|\mathbf{Y}_{\{t,m,N\}}\| \geq \varepsilon \cap \|\mathbf{W}_{\{t,m,N\}}\| \geq \sigma^2\right) \\
&\leq \underbrace{\mathbb{P}\left(\|\mathbf{Y}_{\{t,m,N\}}\| \geq \varepsilon \cap \|\mathbf{W}_{\{t,m,N\}}\| \leq \sigma^2\right)}_{\text{(a)}} + \underbrace{\mathbb{P}\left(\|\mathbf{W}_{\{t,m,N\}}\| \geq \sigma^2\right)}_{\text{(b)}}.
\end{aligned}$$

Using the bound on $\|\mathbf{X}_k\|_2$ from of Step 2, we can directly apply Proposition 3 to bound (a) for any fixed σ^2. To bound the part (b), we see that for any $r \in [1, \ldots, \{m, m, N\}]$ the predictable quadratic variation process takes the form

$$\mathbf{W}_r = \sum_{\{s,e,j\} \leq r} \mathbb{E}\left[\frac{1}{N^2}\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s,e,j}}{\overline{p}_{s,e}}\right)^2 \mathbf{v}_e \mathbf{v}_e^\mathsf{T} \mathbf{v}_e \mathbf{v}_e^\mathsf{T} \,\middle|\, \mathcal{F}_{\{s,e,j\}-1}\right].$$

To show that its norm is large only with a small probability, we use the following lemma, proved later in Section 5.

Lemma 3 (Low probability of the large norm of the predictable quadratic variation process)**.**

$$\mathbb{P}\left(\|\mathbf{W}_{\{t,m,N\}}\| \geq \frac{9\alpha^2 n \log(\kappa n)}{N}\right) \leq n \cdot \exp\left\{-\frac{N}{\alpha^2(n-1)}\right\}$$

Combining Proposition 3 with $\sigma^2 = \frac{9\alpha^2 n \log(\kappa n)}{N}$ and Lemma 3 we obtain

$$\begin{aligned}
&\mathbb{P}\left(\|\widehat{\mathbf{Y}}_{\{t,m,N\}}\| \geq \varepsilon\right) \\
&\quad\leq \mathbb{P}\left(\|\mathbf{Y}_{\{t,m,N\}}\| \geq \varepsilon \cap \|\mathbf{W}_{\{t,m,N\}}\| \leq \sigma^2\right) + \mathbb{P}\left(\|\mathbf{W}_{\{t,m,N\}}\| \geq \sigma^2\right) \\
&\quad\leq 2n \cdot \exp\left\{-\frac{\varepsilon^2/2}{\frac{9\alpha^2 n \log(\kappa n)}{N} + \frac{\varepsilon\alpha^2(n-1)}{3N}}\right\} + n \cdot \exp\left\{-\frac{N}{\alpha^2(n-1)}\right\} \\
&\quad\leq 2n \cdot \exp\left\{-\frac{\varepsilon^2 N}{2(9+\varepsilon/3)\alpha^2 n \log(\kappa n)}\right\} + n \cdot \exp\left\{-\frac{N}{\alpha^2(n-1)}\right\} \\
&\quad\leq 3n \cdot \exp\left\{-\frac{\varepsilon^2 N}{20\alpha^2 n \log(\kappa n)}\right\}.
\end{aligned}$$

Finally, choosing N as in Algorithm 1 we obtain the desired result

$$\begin{aligned}
&\mathbb{P}\left(\|\widehat{\mathbf{Y}}_{\{t,m,N\}}\| \geq \varepsilon\right) \leq 3n \cdot \exp\left\{-\frac{\varepsilon^2 N}{20\alpha^2 n \log(\kappa n)}\right\} = 3n \cdot \exp\left\{-\frac{2\log^2(3\kappa m/\delta)}{\log(\kappa n)}\right\} \\
&\quad\leq 3n \cdot \exp\left\{-2\log(3\kappa m/\delta)\right\} = 3n\left(\frac{\delta}{3\kappa m}\right)^2 \leq \frac{\delta}{3\kappa m} \leq \frac{\delta}{2m}
\end{aligned}$$

5 Proof of Lemma 3 (bound on predictable quadratic variation)

Step 1 (a preliminary bound). We start by rewriting $\mathbf{W}_r$ as

$$\mathbf{W}_r = \frac{1}{N^2} \sum_{\{s,e,j\} \leq r} \mathbb{E}\left[\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} - \frac{z_{s,e,j}}{\overline{p}_{s,e}}\right)^2 \,\middle|\, \mathcal{F}_{\{s,e,j\}-1}\right] \mathbf{v}_e \mathbf{v}_e^\mathsf{T} \mathbf{v}_e \mathbf{v}_e^\mathsf{T}.$$

We rewrite the expectation terms in the equation above as

$$
\begin{aligned}
&\mathbb{E}\left[\left(\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}-\frac{z_{s,e,j}}{\overline{p}_{s,e}}\right)^2\,\middle|\,\mathcal{F}_{\{s,e,j\}-1}\right]\\
&=\mathbb{E}\left[\frac{z^2_{s-1,e,j}}{\overline{p}^2_{s-1,e}}-2\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\frac{z_{s,e,j}}{\overline{p}_{s,e}}+\frac{z^2_{s,e,j}}{\overline{p}^2_{s,e}}\,\middle|\,\mathcal{F}_{\{s,e,j\}-1}\right]\\
&\overset{(a)}{=}\mathbb{E}\left[\frac{z^2_{s-1,e,j}}{\overline{p}^2_{s-1,e}}-2\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\frac{z_{s,e,j}}{\overline{p}_{s,e}}+\frac{z^2_{s,e,j}}{\overline{p}^2_{s,e}}\,\middle|\,\mathcal{F}_{\{s-1,m,N\}}\right]\\
&=\frac{z^2_{s-1,e,j}}{\overline{p}^2_{s-1,e}}-2\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\frac{1}{\overline{p}_{s,e}}\mathbb{E}\left[z_{s,e,j}\,\middle|\,\mathcal{F}_{\{s-1,m,N\}}\right]+\frac{1}{\overline{p}^2_{s,e}}\mathbb{E}\left[z^2_{s,e,j}\,\middle|\,\mathcal{F}_{\{s-1,m,N\}}\right]\\
&\overset{(b)}{=}\frac{z_{s-1,e,j}}{\overline{p}^2_{s-1,e}}-2\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}+\frac{1}{\overline{p}^2_{s,e}}\mathbb{E}\left[z_{s,e,j}\,\middle|\,\mathcal{F}_{\{s-1,m,N\}}\right]\\
&=\frac{1}{\overline{p}^2_{s,e}}\mathbb{E}\left[z_{s,e,j}\,\middle|\,\mathcal{F}_{\{s-1,m,N\}}\right]-\frac{z_{s-1,e,j}}{\overline{p}^2_{s-1,e}}\\
&\overset{(c)}{=}\frac{1}{\overline{p}_{s,e}}\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}-\frac{z_{s-1,e,j}}{\overline{p}^2_{s-1,e}}=\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\left(\frac{1}{\overline{p}_{s,e}}-\frac{1}{\overline{p}_{s-1,e}}\right),
\end{aligned}
$$

where in (a) we use the fact that the approximate probabilities $\overline{p}_{s-1,e}$ and $\overline{p}_{s,e}$ and $z_{s-1,e,j}$ are fixed at the end of the previous iteration, while in (b) and (c) we use the fact that $z_{s,e,j}$ is a Bernoulli of parameter $\overline{p}_{s,e}/\overline{p}_{s-1,e}$ (whenever $z_{s-1,e,j}$ is equal to 1). Therefore we can write $\mathbf{W}$ at the end of any iteration t as

$$
\mathbf{W}_{\{t,m,N\}}=\frac{1}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\sum_{s=1}^{t}\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\left(\frac{1}{\overline{p}_{s,e}}-\frac{1}{\overline{p}_{s-1,e}}\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}.
$$

We can now upper-bound $\mathbf{W}$ as

$$
\begin{aligned}
\mathbf{W}_{\{t,m,N\}}&=\frac{1}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\sum_{s=1}^{t}\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\left(\frac{1}{\overline{p}_{s,e}}-\frac{1}{\overline{p}_{s-1,e}}\right)\\
&\preceq\frac{1}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\left(\max_{s=1\ldots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right)\sum_{s=1}^{t}\left(\frac{1}{\overline{p}_{s,e}}-\frac{1}{\overline{p}_{s-1,e}}\right)\\
&=\frac{1}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\left(\max_{s=1\ldots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right)\left(\frac{1}{\overline{p}_{t,e}}-1\right)\\
&\preceq\frac{1}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}_{t,e}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\left(\max_{s=1\ldots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right),
\end{aligned}
$$

where in the first inequality we use the fact that $\left(\frac{1}{\overline{p}_{s,e}}-\frac{1}{\overline{p}_{s-1,e}}\right)$ and $\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}$ are always positive and the definition of Löwner's ordering and in the second inequality we just upper-bound the last term by $1/\overline{p}_{t,e}$. We further bound the previous term by studying $\overline{p}_{t,e}$. If $\|\mathbf{Y}_{\{t-1,m,N\}}\|\leq\varepsilon$, then $\overline{p}_{t,e}=\widetilde{p}_{t,e}$ and by Proposition 2, $\overline{p}_{t,e}$ is α-accurate and hence

$$
\frac{1}{\overline{p}_{t,e}}=\frac{1}{\widetilde{p}_{t,e}}\leq\frac{\alpha^2}{p_{t,e}}\leq\frac{\alpha^2}{p_{m,e}},
$$

where we use the fact that probabilities (as much as effective resistances) are non-increasing over iterations. Furthermore, in the case if $\|\mathbf{Y}_{\{t-1,m,N\}}\|>\varepsilon$, let $s<t$ be the last time when $\|\mathbf{Y}_{\{s,m,N\}}\|\leq\varepsilon$. In this case, $\overline{p}_{t,e}=\overline{p}_{s,e}$, since the process is "frozen" and therefore,

$$
\frac{1}{\overline{p}_{t,e}}=\frac{1}{\overline{p}_{s,e}}\leq\frac{\alpha^2}{p_{s,e}}\leq\frac{\alpha^2}{p_{m,e}},
$$

since at iteration s, the probability $\overline{p}_{s,e}$ is guaranteed to be α-accurate. As a result, we obtain

$$\begin{aligned}\mathbf{W}_{\{t,m,N\}} &\preceq \frac{1}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}_{t,e}}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\left(\max_{s=1\dots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right)\\ &\preceq \frac{1}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{\alpha^2}{p_{m,e}}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\left(\max_{s=1\dots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right).\end{aligned}$$

Using the definition of $p_{m,e} = a_e r_{m,e}/(n-1)$ and $\mathbf{v}_e = \sqrt{a_e}\mathbf{L}_\mathcal{G}^{-1/2}\mathbf{b}_e$ we have $\mathbf{v}_e^\mathsf{T}\mathbf{v}_e = a_e r_{m,e}$ and thus

$$\begin{aligned}\mathbf{W}_{\{t,m,N\}} &\preceq \frac{\alpha^2}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{p_{m,e}}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\left(\max_{s=1\dots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right)\\ &= \frac{\alpha^2}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{(n-1)a_e r_{m,e}}{a_e r_{m,e}}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\left(\max_{s=1\dots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right)\\ &= \frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\left(\max_{s=1\dots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right). \qquad (5)\end{aligned}$$

Step 2 (introduction of a stochastically dominant process). We want to study

$$\max_{s=1\dots t}\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}} = \max_{s=0\dots t-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}.$$

We know trivially that this quantity is larger or equal than 1 because $z_{0,e,j}/\overline{p}_{0,e} = 1$, but upper-bounding this quantity is not trivial as the evolution of the various $\overline{p}_{s,e}$ depends in a complex way on the interaction between the random variables $z_{s,e,j}$. To simplify the following analysis, we introduce the random variable $z_{t+1,e,j}/\overline{p}_{t+1,e}$, where $\overline{p}_{t+1,e}$ is deterministically defined as $\overline{p}_{t+1,e} = p_{t,e}/\alpha^2$ and $z_{t+1,e,j}$ is defined, coherently with the other $z_{s,e,j}$, as 0 if $z_{t,e,j} = 0$ or as the result of a Bernoulli flip with probability $\overline{p}_{t+1,e}/\overline{p}_{t,e}$ if $z_{t,e,j} = 1$. We have that

$$\max_{s=0\dots t-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}} \le \max_{s=0\dots t}\frac{z_{s,e,j}}{\overline{p}_{s,e}} \le \max_{s=0\dots t+1}\frac{z_{s,e,j}}{\overline{p}_{s,e}},$$

but now we know exactly the value of the last $\overline{p}_{t+1,e}$. Morevoer, we did not lose much since the bounds are tight for the (realizable) worst case when $\overline{p}_{t-1,e} = \overline{p}_{t,e} = \overline{p}_{t+1,e} = p_{t,e}$. Finding a deterministic bound on $\max_{s=0\dots t+1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}$ is not easy, since $\overline{p}_{t+1,e} = p_{t,e}$ can be very close to 0. Nonetheless, whenever $\overline{p}_{s,e}$ is significantly smaller than $\overline{p}_{s-1,e}$, the probability of keeping a copy of edge e at iteration s (i.e., $z_{s,e,j} = 1$) is also very small. As a result, we expect the ratio $\frac{z_{s,e,j}}{\overline{p}_{s,e}}$ to be still small with high probability. Unfortunately, due to the dependency between different copies of the edges at different iterations, it seems difficult to exploit this intuition directly to provide an overall high-probability bound on $\mathbf{W}_{\{t,m,N\}}$. For this reason, we simplify the analysis by replacing each of the (potentially dependent) chains $\{z_{s,e,j}/\overline{p}_{s,e}\}_{s=0}^{t+1}$ with a set of (independent) random variables $w_{0,e,j}$ that will stochastically dominate them.

We define the random variable $w_{s,e,j}$ using the following conditional distribution[3]

$$\mathbb{P}\left(\frac{1}{w_{s,e,j}} \le a \;\middle|\; \mathcal{F}_{\{s,m,N\}}\right) = \begin{cases}0 & \text{for} \quad a < 1/\overline{p}_{s,e}\\ 1-\frac{1}{\overline{p}_{s,e}a} & \text{for} \quad 1/\overline{p}_{s,e} \le a < \alpha^2/p_{t,e}\\ 1 & \text{for} \quad \alpha^2/p_{t,e} \le a\end{cases}.$$

[3] Notice that unlike $z_{s,e,j}$, $w_{s,e,j}$ is no longer $\mathcal{F}_{\{s,m,N\}}$-measurable but it is $\mathcal{F}'_{\{s,m,N\}}$-measurable, where

$$\mathcal{F}'_{\{s,e,j\}} = \left\{u_{s',e',j'},\ \forall\{s',e',j'\} \le \{s,e,j\}\right\} \cup \{w_{s,e,j}\} = \mathcal{F}_{\{s,e,j\}} \cup \{w_{s,e,j}\}.$$

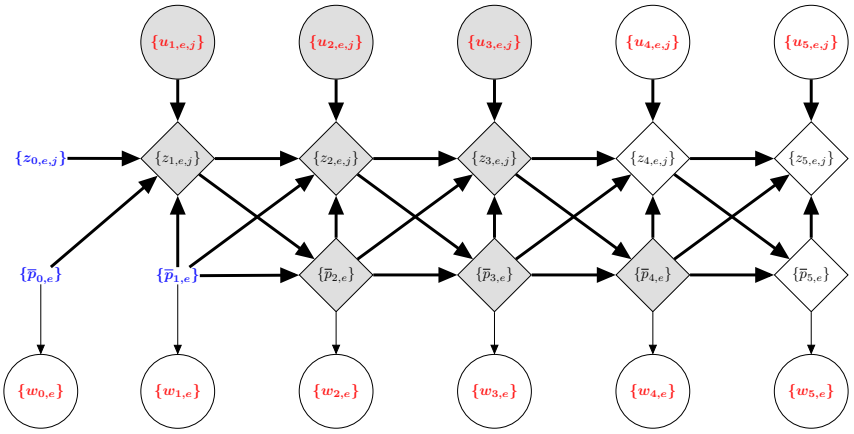

Figure 1: The dependence graph of the considered variables. **Red** variables are **random**, Black variables are deterministically computed using their input (a function of their input), with bold lines indicating the deterministic (functional) relation. **Blue** variables are **constants**. A grey filling indicates that a random variable is observed, or a function of observed variables.

Note that the distribution of $\frac{1}{w_{s,e,j}}$ conditioned on $\mathcal{F}_{\{s,m,N\}}$ is determined by only $\overline{p}_{s,e}$, $p_{t,e}$, and α, where $p_{t,e}$ and α are fixed. Remembering that $\overline{p}_{s,e}$ is a function of $\mathcal{F}_{\{s-1,m,N\}}$ (computed using the previous iteration), we have that

$$\mathbb{P}\left(\frac{1}{w_{s,e,j}} \leq a \,\middle|\, \mathcal{F}_{\{s,m,N\}}\right) = \mathbb{P}\left(\frac{1}{w_{s,e,j}} \leq a \,\middle|\, \mathcal{F}_{\{s-1,m,N\}}\right).$$

Notice that in the definition of $w_{s,e,j}$, none of the other $w_{s',e',j'}$ (for any different s', e', or j') appears and $\overline{p}_{s,e}$ is a function of $\mathcal{F}_{\{s-1,m,N\}}$. It follows that given $\mathcal{F}_{\{s-1,m,N\}}$, $w_{s,e,j}$ is independent from all other $w_{s',e',j'}$ (for any different s', e', or j'). This is easier to see in the probabilistic graphical model reported in Figure 1, which illustrates the dependence between the various variables.

Finally for the special case $w_{0,e,j}$ the definition above reduces to

$$\mathbb{P}\left(\frac{1}{w_{0,e,j}} \leq a\right) = \begin{cases} 0 & \text{for} \quad a < 1 \\ 1 - \frac{1}{a} & \text{for} \quad 1 \leq a < \alpha^2/p_{t,e} \\ 1 & \text{for} \quad \alpha^2/p_{t,e} \leq a \end{cases}, \tag{6}$$

since $\overline{p}_{0,e} = 1$ by definition. From this definition, $w_{0,e,j}$ and $w_{0,e',j'}$ are all independent, and this will allow us to use stronger concentration inequalities for independent random variables.

We remind the reader that a random variable A stochastically dominates random variable B, if for all values a the two equivalent conditions are verified

$$\mathbb{P}(A \geq a) \geq \mathbb{P}(B \geq a) \Leftrightarrow \mathbb{P}(A \leq a) \leq \mathbb{P}(B \leq a).$$

As a consequence, if A dominates B the following implication holds

$$\mathbb{P}(A \geq a) \geq \mathbb{P}(B \geq a) \implies \mathbb{E}[A] \geq \mathbb{E}[B],$$

while the reverse (A dominates B if $\mathbb{E}[A] \geq \mathbb{E}[B]$) is not true in general. Following this definition of stochastic dominance, our goal is to prove

$$\mathbb{P}\left(\max_{s=0\ldots t+1} \frac{z_{s,e,j}}{\overline{p}_{s,e}} \leq a\right) \geq \mathbb{P}\left(\frac{1}{w_{0,e,j}} \leq a\right).$$

We prove this inequality by proceeding backward with a sequence of conditional probabilities.

Step 3 (case $t+1$). We first study the distribution of the maximum conditional to the state of the algorithm at the end of iteration t, i.e., $\mathcal{F}_{\{t,m,N\}}$, that is

$$\rho_a \stackrel{\text{def}}{=\!=} \mathbb{P}\left(\max\left\{\max_{s=0\ldots t} \frac{z_{s,e,j}}{\overline{p}_{s,e}}; \frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\} \leq a \,\middle|\, \mathcal{F}_{\{t,m,N\}}\right).$$

We distinguish two cases, depending on whether $z_{t,e,j} = 0$ or 1, which directly depends on $\mathcal{F}_{\{t,m,N\}}$.

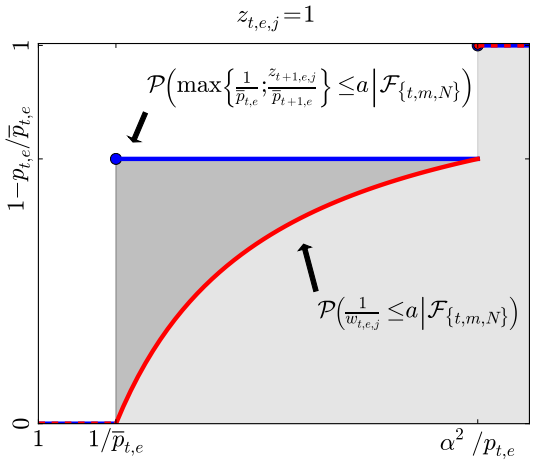

Figure 2: C.d.f. of $\max\left\{1/\overline{p}_{t,e};\, z_{t+1,e,j}/\overline{p}_{t+1,e}\right\}$ and $1/w_{t,e,j}$ conditioned on $\mathcal{F}_{\{t,m,N\}}$

- If $z_{t,e,j}=0$, then by definition $z_{t+1,e,j}=0$ and thus for any a

$$\rho_a=\mathbb{P}\left(\max_{s=0\ldots t-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right),$$

  since the maximum is always attained by one of the terms for which $z_{s,e,j}=1$.

- If $z_{t,e,j}=1$, then by definition $z_{t+1,e,j}\big|\mathcal{F}_{\{t,m,N\}}$ is a Bernoulli variable of parameter $\overline{p}_{e,t+1}/\overline{p}_{e,t}$. Since $z_{t,e,j}=1$ and $\overline{p}_{t,e}\leq\overline{p}_{s,e}$ for all $s\leq t$, the inner maximum in ρ_a is always attained by $\frac{z_{t,e,j}}{\overline{p}_{t,e}}$ and thus we can write ρ_a as

$$\rho_a=\mathbb{P}\left(\max\left\{\frac{1}{\overline{p}_{t,e}};\frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right).$$

  We can easily study the c.d.f. of the maximum depending on the value of a as follows:

  - If $a<1/\overline{p}_{t,e}$, then $\rho_a=0$, since the maximum is at least $\frac{1}{\overline{p}_{t,e}}$.
  - If $1/\overline{p}_{t,e}\leq a<1/\overline{p}_{t+1,e}$, we are in the case where the Bernoulli variable takes value 0, and thus

$$\rho_a=\mathbb{P}\left(z_{t+1,e,j}=0\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right)=1-\frac{\overline{p}_{t+1,e}}{\overline{p}_{t,e}}.$$

  - If $a\geq 1/\overline{p}_{t+1,e}$, then $\rho_a=1$, since a is bigger or equal than the largest value attainable by the maximum.

We now move to analyze the distribution

$$\pi_a\overset{\text{def}}{=\!=}\mathbb{P}\left(\max\left\{\max_{s=0\ldots t}\frac{z_{s,e,j}}{\overline{p}_{s,e}};\frac{z_{t,e,j}}{w_{t,e,j}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right),$$

where we replaced $z_{t+1,e,j}/\overline{p}_{t+1,e}$ with the variable $z_{t,e,j}/w_{t,e,j}$. Note that in this expression only $w_{t,e,j}$ is random, and all other terms are fixed given $\mathcal{F}_{\{t,m,N\}}$. If $z_{t,e,j}=0$, then for any a,

$$\pi_a=\mathbb{P}\left(\max_{s=0\ldots t-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right).$$

On the other hand, if $z_{t,e,j}=1$, we can use the definition of $w_{t,e,j}$ to determine its c.d.f.. As a result, we notice that if $z_{t,e,j}=0$, then $\rho_a=\pi_a$ for any a, while for $z_{t,e,j}=1$ we have

$$
\begin{aligned}
&\mathbb{P}\left(\max\left\{\max_{s=0\dots t}\frac{z_{s,e,j}}{\overline{p}_{t,e}};\frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right)\\
&=\begin{cases}0 & \text{for}\quad 1/\overline{p}_{t,e}<a \text{ and } z_{t,e,j}=1\\ 1-\frac{\overline{p}_{t+1,e}}{\overline{p}_{t,e}} & \text{for}\quad 1/\overline{p}_{t,e}\leq a<1/\overline{p}_{t+1,e} \text{ and } z_{t,e,j}=1\\ 1 & \text{for}\quad 1/\overline{p}_{t+1,e}\leq a \text{ and } z_{t,e,j}=1\end{cases}\\
&\geq\begin{cases}0 & \text{for}\quad 1/\overline{p}_{t,e}<a \text{ and } z_{t,e,j}=1\\ 1-\frac{1}{\overline{p}_{t,e}a} & \text{for}\quad 1/\overline{p}_{t,e}\leq a<1/\overline{p}_{t+1,e} \text{ and } z_{t,e,j}=1\\ 1 & \text{for}\quad 1/\overline{p}_{t+1,e}=\alpha^2/p_{t,e}\leq a \text{ and } z_{t,e,j}=1\end{cases}\\
&=\mathbb{P}\left(\max\left\{\max_{s=0\dots t}\frac{z_{s,e,j}}{\overline{p}_{t,e}};\frac{z_{t,e,j}}{w_{t,e,j}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right),
\end{aligned}
$$

where we used the definition $\overline{p}_{t+1,e}=p_{t,e}/\alpha^2$, and the fact that since $\overline{p}_{t,e}$ is always either α-accurate, or frozen to a previous α-accurate value $\widetilde{p}_{s,e}$, we have that $\overline{p}_{t,e}\geq\overline{p}_{t+1,e}=p_{t,e}/\alpha^2$. The inequality between ρ_a and π_a comes from the fact that in the intermediate interval a is such that $a<1/\overline{p}_{t+1,e}$ and then $1-\frac{1}{\overline{p}_{t,e}a}\leq 1-\frac{\overline{p}_{t+1,e}}{\overline{p}_{t,e}}$, as it is easy to observe in Fig. 2. We can thus conclude that $\rho_a\geq\pi_a$ for any a.

Now that we introduced $w_{t,e,j}$ in the maximum of the conditional probability, we can focus on the overall probability as

$$
\begin{aligned}
&\mathbb{P}\left(\max_{s=0\dots t+1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}\leq a\right)=\sum_{\mathcal{F}_{\{t,m,N\}}}\mathbb{P}\left(\max_{s=0\dots t+1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right)\mathbb{P}(\mathcal{F}_{\{t,m,N\}})\\
&=\sum_{\mathcal{F}_{\{t,m,N\}}}\mathbb{P}\left(\max\left\{\max_{s=0\dots t}\frac{z_{s,e,j}}{\overline{p}_{s,e}};\frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right)\mathbb{P}(\mathcal{F}_{\{t,m,N\}})\\
&\geq\sum_{\mathcal{F}_{\{t,m,N\}}}\mathbb{P}\left(\max\left\{\max_{s=0\dots t}\frac{z_{s,e,j}}{\overline{p}_{s,e}};\frac{z_{t,e,j}}{w_{t,e,j}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{t,m,N\}}\right)\mathbb{P}(\mathcal{F}_{\{t,m,N\}})\\
&=\mathbb{P}\left(\max\left\{\max_{s=0\dots t}\frac{z_{s,e,j}}{\overline{p}_{s,e}};\frac{z_{t,e,j}}{w_{t,e,j}}\right\}\leq a\right).
\end{aligned}
$$

Step 4 (generic case). We now proceed by peeling off layers from the end of the chain one by one. We show how to move from an iteration $k\leq t$ to $k-1$.

$$
\begin{aligned}
&\mathbb{P}\left(\max\left\{\max_{s=0\dots k}\frac{z_{s,e,j}}{\overline{p}_{s,e}};\frac{z_{k,e,j}}{w_{k,e,j}}\right\}\leq a\right)=\mathbb{P}\left(\max\left\{\max_{s=0\dots k-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}};\frac{z_{k,e,j}}{\overline{p}_{k,e}};\frac{z_{k,e,j}}{w_{k,e,j}}\right\}\leq a\right) \qquad (7)\\
&=\sum_{\mathcal{F}_{\{k-1,m,N\}}}\mathbb{P}\left(\max\left\{\max_{s=0\dots k-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}};\frac{z_{k,e,j}}{\overline{p}_{k,e}};\frac{z_{k,e,j}}{w_{k,e,j}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{k-1,m,N\}}\right)\mathbb{P}(\mathcal{F}_{\{k-1,m,N\}})\\
&=\sum_{\mathcal{F}_{\{k-1,m,N\}}}\mathbb{P}\left(\max\left\{\max_{s=0\dots k-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}};z_{k,e,j}\max\left\{\frac{1}{\overline{p}_{k,e}};\frac{1}{w_{k,e,j}}\right\}\right\}\leq a\;\middle|\;\mathcal{F}_{\{k-1,m,N\}}\right)\mathbb{P}(\mathcal{F}_{\{k-1,m,N\}})\\
&=\sum_{\mathcal{F}_{\{k-1,m,N\}}}\mathbb{P}\left(\max\left\{\max_{s=0\dots k-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}};z_{k-1,e,j}\mathbb{I}\left\{u_{k,e,j}\leq\frac{\overline{p}_{k,e}}{\overline{p}_{k-1,e}}\right\}\frac{1}{w_{k,e,j}}\right\}\leq a\;\middle|\;\mathcal{F}_{\{k-1,m,N\}}\right)\mathbb{P}(\mathcal{F}_{\{k-1,m,N\}}),
\end{aligned}
$$

where in the last equality we use the fact that the inner maximum is always attained by $1/w_{k,e,j}$ since by definition $1/w_{k,e,j}$ is lower-bounded by $1/\overline{p}_{k,e}$ and we use the recursive definition of $z_{k,e,j}$.

Since $u_{k,e,j}$ and $w_{k,e,j}$ are independent given $\mathcal{F}_{\{k-1,m,N\}}$, we have

$$\begin{aligned}
&\mathbb{P}\left(\mathbb{I}\left\{u_{k,e,j} \le \frac{\overline{p}_{k,e}}{\overline{p}_{k-1,e}}\right\}\frac{1}{w_{k,e,j}} \le a \;\middle|\; \mathcal{F}_{\{k-1,m,N\}}\right) \\
&= \begin{cases} 0 & \text{for} \quad a \le 0 \\ 1-\frac{\overline{p}_{k,e}}{\overline{p}_{k-1,e}} & \text{for} \quad 0 \le a < 1/\overline{p}_{k,e} \\ 1-\frac{\overline{p}_{k,e}}{\overline{p}_{k-1,e}} + \frac{\overline{p}_{k,e}}{\overline{p}_{k-1,e}}\left(1-\frac{1}{\overline{p}_{k,e}a}\right) = 1-\frac{1}{\overline{p}_{k-1,e}a} & \text{for} \quad 1/\overline{p}_{k,e} \le a < \alpha^2/p_{t,e} \\ 1 & \text{for} \quad \alpha^2/p_{t,e} \le a \end{cases} \\
&\ge \begin{cases} 0 & \text{for} \quad a < 1/\overline{p}_{k-1,e} \\ 1-\frac{1}{\overline{p}_{k-1,e}a} & \text{for} \quad 1/\overline{p}_{k-1,e} \le a < 1/\overline{p}_{k,e} \\ 1-\frac{1}{\overline{p}_{k-1,e}a} & \text{for} \quad 1/\overline{p}_{k,e} \le a < \alpha^2/p_{t,e} \\ 1 & \text{for} \quad \alpha^2/p_{t,e} \le a \end{cases} \\
&= \mathbb{P}\left(\frac{1}{w_{k-1,e,j}} \le a \;\middle|\; \mathcal{F}_{\{k-2,m,N\}}\right) = \mathbb{P}\left(\frac{1}{w_{k-1,e,j}} \le a \;\middle|\; \mathcal{F}_{\{k-1,m,N\}}\right).
\end{aligned} \tag{8}$$

Since given $\mathcal{F}_{\{k-1,m,N\}}$, everything in (7) is fixed except $u_{k,e,j}$ and $w_{k,e,j}$, we can use the stochastic dominance in (8) to obtain

$$\mathbb{P}\left(\max\left\{\max_{s=0\dots k}\frac{z_{s,e,j}}{\overline{p}_{s,e}}; \frac{z_{k,e,j}}{w_{k,e,j}}\right\} \le a\right) \ge \mathbb{P}\left(\max\left\{\max_{s=0\dots k-1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}; \frac{z_{k-1,e,j}}{w_{k-1,e,j}}\right\} \le a\right).$$

Applying the inequality recursively from $k=t$ to $k=1$ removes all $z_{s,e,j}$ from the maximum and we are finally left with only $w_{0,e,j}$ as we wanted,

$$\mathbb{P}\left(\max_{s=0\dots t+1}\frac{z_{s,e,j}}{\overline{p}_{s,e}} \le a\right) \ge \mathbb{P}\left(\max\left\{\frac{z_{0,e,j}}{\overline{p}_{0,e}}; \frac{z_{0,e,j}}{w_{0,e,j}}\right\} \le a\right) \ge \mathbb{P}\left(\frac{1}{w_{0,e,j}} \le a\right),$$

where in the last inequality we used that $z_{0,e,j}=1$ from the definition of the algorithm and $\overline{p}_{0,e}=1$ while $w_{0,e,j} \le 1$ by (6).

Step 5 (stochastic dominance on $\mathbf{W}_{\{t,m,N\}}$). Now that we proved the stochastic dominance of $1/w_{0,e,j}$, we plug this result in the definition of $\mathbf{W}_{\{t,m,N\}}$. For the sake of notation, we introduce the term $\overline{p}^{\max}_{t+1,e,j}$ to indicate the maximum over the first $t+1$ step of copy e,j such that

$$\max_{s=0\dots t+1}\frac{z_{s,e,j}}{\overline{p}_{s,e}} = \frac{1}{\overline{p}^{\max}_{t+1,e,j}}.$$

We first notice that while $\mathbf{Y}_{\{t,m,N\}}$ is not necessarily PSD, $\mathbf{W}_{\{t,m,N\}}$ is a sum of PSD matrices. Introducing the function $\Lambda(\{1/\overline{p}^{\max}_{t+1,e,j}\}_{e,j})$ we can restate Eq. 5 as

$$\|\mathbf{W}_{\{t,m,N\}}\| = \lambda_{\max}(\mathbf{W}_{\{t,m,N\}}) \le \Lambda(\{1/\overline{p}^{\max}_{t+1,e,j}\}_{e,j}) \overset{\text{def}}{=} \lambda_{\max}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}^{\max}_{t+1,e,j}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right).$$

In Step 4, we showed that $1/\overline{p}^{\max}_{t+1,e,j}$ is stochastically dominated by $1/w_{0,e,j}$ for every e and j. In order to bound $\Lambda(\{1/\overline{p}^{\max}_{t+1,e,j}\}_{e,j})$ we need to show that this dominance also applies to the summation over all edges inside the matrix norm. Since the matrix $\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}^{\max}_{t+1,e,j}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}$ is symmetric, we can reformulate $\Lambda(\{1/\overline{p}^{\max}_{t+1,e,j}\}_{e,j})$ as

$$\begin{aligned}
&\lambda_{\max}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}^{\max}_{t+1,e,j}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right) = \max_{\mathbf{x}:\|\mathbf{x}\|=1}\mathbf{x}^{\mathsf{T}}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}^{\max}_{t+1,e,j}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right)\mathbf{x} \\
&= \max_{\mathbf{x}:\|\mathbf{x}\|=1}\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}^{\max}_{t+1,e,j}}\mathbf{x}^{\mathsf{T}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\mathbf{x} = \max_{\mathbf{x}:\|\mathbf{x}\|=1}\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{\overline{p}^{\max}_{t+1,e,j}}\left(\mathbf{v}_e^{\mathsf{T}}\mathbf{x}\right)^2.
\end{aligned}$$

From this reformulation, it is easy to see that the function $\Lambda(\{1/\overline{p}^{\max}_{t+1,e,j}\}_{e,j})$ is monotonically increasing w.r.t. the individual $1/\overline{p}^{\max}_{t+1,e,j}$, or in other words that increasing an $1/\overline{p}^{\max}_{t+1,e,j}$ without decreasing the others can only increase the maximum. Introducing $\Lambda(\{1/w_{0,e,j}\}_{e,j})$ as

$$\Lambda(\{1/w_{0,e,j}\}_{e,j}) \stackrel{\text{def}}{=} \max_{\mathbf{x}:\|\mathbf{x}\|=1} \frac{\alpha^2(n-1)}{N^2} \sum_{j=1}^{N}\sum_{e=1}^{m} \frac{1}{w_{0,e,j}} \left(\mathbf{v}_e^\mathsf{T}\mathbf{x}\right)^2 .$$

we now need to prove the stochastic dominance of $\Lambda(\{1/w_{0,e,j}\}_{e,j})$ over $\Lambda(\{1/\overline{p}^{\max}_{t+1,e,j}\}_{e,j})$. Using the definition of $1/\overline{p}^{\max}_{t+1,e,j}$ we have

$$\begin{aligned}
&\mathbb{P}\left(\Lambda\left(\left\{\frac{1}{\overline{p}^{\max}_{t+1,e,j}}\right\}_{e,j}\right) \le a\right) = \mathbb{P}\left(\Lambda\left(\left\{\max_{s=0\dots t+1}\frac{z_{s,e,j}}{\overline{p}_{s,e}}\right\}_{e,j}\right) \le a\right) \\
&= \mathbb{P}\left(\Lambda\left(\left\{\max\left\{\max_{s=0\dots t}\frac{z_{s,e,j}}{\overline{p}_{s,e}}; \frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\}\right\}_{e,j}\right) \le a\right) \\
&= \sum_{\mathcal{F}_{\{t,m,N\}}} \mathbb{P}\left(\Lambda\left(\left\{\max\left\{\frac{1}{\overline{p}^{\max}_{t,e,j}}; \frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\}\right\}_{e,j}\right) \le a \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right) \mathbb{P}\left(\mathcal{F}_{\{t,m,N\}}\right).
\end{aligned}$$

For a fixed $\mathcal{F}_{\{t,m,N\}}$, denote with $A_{e,j}$ the random variable $\max\left\{\frac{1}{\overline{p}^{\max}_{t,e,j}}; \frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\}$, and with A the set of random variables $\{A_{e,j}\}_{e,j}$ for all e and j. Similarly, we define $B_{e,j}$ and B for the random variable $\max\left\{\frac{1}{\overline{p}^{\max}_{t,e,j}}; \frac{z_{t,e,j}}{w_{t,e,j}}\right\}$. Note that, given $\mathcal{F}_{\{t,m,N\}}$, $\frac{1}{\overline{p}^{\max}_{t,e,j}}$ is a constant and $\max\left\{\frac{1}{\overline{p}^{\max}_{t,e,j}}; x\right\}$ is a monotonically increasing function in x. Therefore, given $\mathcal{F}_{\{t,m,N\}}$, $\frac{z_{t,e,j}}{w_{t,e,j}}$ stochastically dominates $\frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}$ and $B_{e,j}$ dominates $A_{e,j}$, since stochastic dominance is preserved by monotone functions [3]. We have

$$\begin{aligned}
&\mathbb{P}\left(\Lambda\left(\left\{\max\left\{\frac{1}{\overline{p}^{\max}_{t,e,j}}; \frac{z_{t+1,e,j}}{\overline{p}_{t+1,e}}\right\}\right\}_{e,j}\right) \le a \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right) \\
&= \mathbb{P}\left(\Lambda\left(A\right) \le a \mid \mathcal{F}_{\{t,m,N\}}\right) = \mathop{\mathbb{E}}_{A}\left[\mathbb{I}\left\{\Lambda\left(A\right) \le a\right\} \mid \mathcal{F}_{\{t,m,N\}}\right] \\
&= \mathop{\mathbb{E}}_{A\setminus A_{k,l}}\left[\mathop{\mathbb{E}}_{A_{k,l}}\left[\mathbb{I}\left\{\Lambda\left(\{A_{e,j}\}_{e,j}\right) \le a\right\} \;\middle|\; \mathcal{F}_{\{t,m,N\}} \cap A\setminus A_{k,l}\right] \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right] \\
&\stackrel{(a)}{\ge} \mathop{\mathbb{E}}_{A\setminus A_{k,l}}\left[\mathop{\mathbb{E}}_{B_{k,l}}\left[\mathbb{I}\left\{\Lambda\left(\{A_{e,j}\}_{(e,j)\neq(k,l)}, B_{k,l}\right) \le a\right\} \;\middle|\; \mathcal{F}_{\{t,m,N\}} \cap A\setminus A_{k,l}\right] \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right] \\
&\stackrel{(b)}{=} \mathop{\mathbb{E}}_{B_{k,l}}\left[\mathop{\mathbb{E}}_{A\setminus A_{k,l}}\left[\mathbb{I}\left\{\Lambda\left(\{A_{e,j}\}_{(e,j)\neq(k,l)}, B_{k,l}\right) \le a\right\} \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right] \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right] \\
&\stackrel{(c)}{\ge} \mathop{\mathbb{E}}_{B}\left[\mathbb{I}\left\{\Lambda\left(\{B_{e,j}\}_{e,j}\right) \le a\right\} \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right] \\
&= \mathbb{P}\left(\Lambda\left(\left\{\max\left\{\frac{1}{\overline{p}^{\max}_{t,e,j}}; \frac{z_{t,e,j}}{w_{t,e,j}}\right\}\right\}_{e,j}\right) \le a \;\middle|\; \mathcal{F}_{\{t,m,N\}}\right)
\end{aligned}$$

where inequality (a) follows from the stochastic dominance and the fact that Λ is monotonically increasing, equality (b) comes from the independence of $B_{e,j}$ from all A given $\mathcal{F}_{\{t,m,N\}}$, and inequality (c) is obtained by repeatedly applying (a) and (b) to replace all $A_{e,j}$ variables with $B_{e,j}$. We can further iterate this inequalities (similarly to Step 4) to obtain the desired result

$$\begin{aligned}
\mathbb{P}(\|\mathbf{W}_{\{t,m,N\}}\| \ge \sigma^2) &\le \mathbb{P}\left(\lambda_{\max}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\left(\max_{s=1\dots t}\left\{\frac{z_{s-1,e,j}}{\overline{p}_{s-1,e}}\right\}\right)\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\right) \ge \sigma^2\right) \\
&\le \mathbb{P}\left(\lambda_{\max}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{w_{0,e,j}}\mathbf{v}_e\mathbf{v}_e^\mathsf{T}\right) \ge \sigma^2\right).
\end{aligned}$$

Step 6 (concentration inequality). Since all $w_{0,e,j}$ are (unconditionally) independent from each other, we can apply the following theorem.

Proposition 4 (Theorem 5.1.1 [7])**.** *Consider a finite sequence $\{\mathbf{X}_k : k = 1, 2, 3, \dots\}$ whose values are independent, random, PSD Hermitian matrices with dimension d. Assume that each term in the sequence is uniformly bounded in the sense that*

$$\lambda_{\max}(\mathbf{X}_k) \leq L \quad \textit{almost surely} \quad \textit{for } k = 1, 2, 3, \dots.$$

Introduce the random matrix $\mathbf{V} \stackrel{\text{def}}{=} \sum_k \mathbf{X}_k$, and the maximum eigenvalue of its expectation

$$\mu_{\max} \stackrel{\text{def}}{=} \lambda_{\max}(\mathbb{E}[\mathbf{V}]) = \lambda_{\max}\left(\sum_k \mathbb{E}[\mathbf{X}_k]\right).$$

Then, for all $h \geq 0$,

$$\begin{aligned}\mathbb{P}(\lambda_{\max}(\mathbf{V}) \geq (1+h)\mu_{\max}) &\leq d \cdot \left[\frac{e^h}{(1+h)^{1+h}}\right]^{\frac{\mu_{\max}}{R}} \\ &\leq d \cdot \exp\left\{-\frac{\mu_{\max}}{R}((h+1)\log(h+1) - h)\right\}.\end{aligned}$$

In our case, we have

$$\mathbf{X}_{\{e,j\}} = \frac{\alpha^2(n-1)}{N^2}\frac{1}{w_{0,e,j}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}} \preceq \frac{\alpha^2(n-1)}{N^2}\frac{\alpha^2}{p_{t,e}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}} \preceq \frac{\alpha^4(n-1)^2 a_e r_{m,e}}{a_e r_{m,e} N^2}\mathbf{I},$$

where the first inequality follows from (6) and the second from the fact that $\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}} \preceq \|\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\|\mathbf{I} = \|\mathbf{v}_e^{\mathsf{T}}\mathbf{v}_e\|\mathbf{I} = a_e r_{m,e}\mathbf{I}$.

Therefore, we can use $L \stackrel{\text{def}}{=} \alpha^4(n-1)^2/N^2$ for the purpose of Proposition 4. We need now to compute $\mathbb{E}[\mathbf{X}_k]$, that we can use in turn to compute $\mu_{\max}$. We begin by computing the expected value of $1/w_{0,e,j}$. Let as denote the c.d.f. of $1/w_{0,e,j}$ as

$$F_{1/w_{0,e,j}}(a) = \mathbb{P}\left(\frac{1}{w_{0,e,j}} \leq a\right).$$

Since $\mathbb{P}(1/w_{0,e,j} \geq 0) = 1$ by (6), we have that

$$\begin{aligned}&\mathbb{E}\left[\frac{1}{w_{0,e,j}}\right] = \int_{a=0}^{\infty}\left[1 - F_{1/w_{0,e,j}}(a)\right]\mathrm{d}a \\ &= \int_{a=0}^{1}\left(1 - F_{1/w_{0,e,j}}(a)\right)\mathrm{d}a + \int_{a=1}^{\alpha^2/p_{t,e}}\left(1 - F_{1/w_{0,e,j}}(a)\right)\mathrm{d}a + \int_{a=\alpha^2/p_{t,e}}^{\infty}\left(1 - F_{1/w_{0,e,j}}(a)\right)\mathrm{d}a \\ &= \int_{a=0}^{1}(1-0)\,\mathrm{d}a + \int_{a=1}^{\alpha^2/p_{t,e}}\left(1 - \left(1 - \frac{1}{a}\right)\right)\mathrm{d}a + \int_{a=\alpha^2/p_{t,e}}^{\infty}(1-1)\,\mathrm{d}a \\ &= \int_{a=0}^{1}\mathrm{d}a + \int_{a=1}^{\alpha^2/p_{t,e}}\frac{1}{a}\,\mathrm{d}a = 1 + \log(\alpha^2/p_{t,e}),\end{aligned}$$

where we used the definition of the c.d.f. of $1/w_{0,e,j}$ in Eq. 6. Thus we have

$$\mathbb{E}\left[\mathbf{X}_{\{e,j\}}\right] = \frac{\alpha^2(n-1)}{N^2}\mathbb{E}\left[\frac{1}{w_{0,e,j}}\right]\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}} = \frac{\alpha^2(n-1)}{N^2}\left(1 + \log(\alpha^2/p_{t,e})\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}.$$

Therefore,

$$\mu_{\max} = \lambda_{\max}(\mathbb{E}[\mathbf{V}]) = \lambda_{\max}\Big(\sum_{\{e,j\}}\mathbb{E}\left[\mathbf{X}_{\{e,j\}}\right]\Big) = \lambda_{\max}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\left(1 + \log(\alpha^2/p_{t,e})\right)\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right).$$

We now find more explicit upper and lower bounds on $\mu_{\max}$. We know that

$$p_{t,e} \geq p_{m,e} = \frac{a_e r_{m,e}}{n-1} = \frac{a_e \mathbf{b}_e^{\mathsf{T}}\mathbf{L}_{\mathcal{G}}^{+}\mathbf{b}_e^{\mathsf{T}}}{n-1} \geq \frac{a_e \lambda_{\min}(\mathbf{L}_{\mathcal{G}}^{+})\|\mathbf{b}_e\|^2}{n-1} = \frac{2a_e}{\lambda_{\max}(\mathbf{L}_{\mathcal{G}})(n-1)} \geq \frac{2a_{\min}}{a_{\max}n(n-1)},$$

where $\lambda_{\min}(\mathbf{L}_{\mathcal{G}}^{+}) = 1/\lambda_{\max}(\mathbf{L}_{\mathcal{G}})$ is the smallest non-zero eigenvalue of $\mathbf{L}_{\mathcal{G}}^{+}$. To show that $\lambda_{\max}(\mathbf{L}_{\mathcal{G}}) \leq a_{\max} n$, denote with $\mathcal{K}_n$ the unweighted complete graph over n edges. We use the fact that

$$\begin{aligned}\lambda_{\max}(\mathbf{L}_{\mathcal{G}}) &= \max_{\|\mathbf{x}\|=1} \mathbf{x}^{\mathsf{T}}\mathbf{L}_{\mathcal{G}}\mathbf{x} = \max_{\|\mathbf{x}\|=1} \frac{1}{2}\sum_{e=1}^{m} a_e (x_{e_i} - x_{e_j})^2 \\ &\leq \max_{\|\mathbf{x}\|=1} \frac{a_{\max}}{2}\sum_{e=1}^{m}(x_{e_i} - x_{e_j})^2 \leq a_{\max} \max_{\|\mathbf{x}\|=1} \frac{1}{2}\sum_{e=1}^{n^2}(x_{e_i} - x_{e_j})^2 \\ &= a_{\max}\lambda_{\max}(\mathbf{L}_{\mathcal{K}_n}) \leq a_{\max} n\end{aligned}$$

where in the first inequality bounds summation over the weighted graph $\mathcal{G}$ with a summation over an unweighted version of $\mathcal{G}$ times $a_{\max}$, in the second inequality we passed from a summation over the graph $\mathcal{G}$ to a summation over the full graph $\mathcal{K}_n$ and in the last passage we bounded the largest eigenvalue of the $\mathbf{L}_{\mathcal{K}_n}$. Using the definition $\kappa^2 = a_{\max}/a_{\min}$ and the assumption $\alpha \leq \sqrt{\kappa n/3}$, then we obtain

$$\begin{aligned}\mu_{\max} &\leq \lambda_{\max}\left(\frac{\alpha^2(n-1)}{N}\sum_{e=1}^{m}(1+\log(\kappa^2 n^2\alpha^2/2))\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right) \\ &\leq \frac{a_{\max}(1+2\log(\kappa n)+2\log(\alpha))\alpha^2(n-1)}{N}\lambda_{\max}\left(\sum_{e=1}^{m}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right) \\ &\leq \frac{3\alpha^2(n-1)\log(\kappa n)}{N}\lambda_{\max}\left(\mathbf{P}\right) \leq \frac{3\alpha^2(n-1)\log(\kappa n)}{N}.\end{aligned}$$

Furthermore, we have that

$$\mu_{\max} \geq \lambda_{\max}\left(\frac{\alpha^2(n-1)}{N}\sum_{e=1}^{m}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right) = \frac{\alpha^2(n-1)}{N}\lambda_{\max}\left(\mathbf{P}\right) \geq \frac{\alpha^2(n-1)}{N}.$$

Therefore, selecting $h = 2$ and applying Proposition 4 we have

$$\begin{aligned}&\mathbb{P}\left(\lambda_{\max}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{w_{0,e,j}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right) \geq \frac{9\alpha^2(n-1)\log(\kappa n)}{N}\right) \\ &\leq \mathbb{P}\left(\lambda_{\max}\left(\frac{\alpha^2(n-1)}{N^2}\sum_{j=1}^{N}\sum_{e=1}^{m}\frac{1}{w_{0,e,j}}\mathbf{v}_e\mathbf{v}_e^{\mathsf{T}}\right) \geq (1+2)\mu_{\max}\right) \\ &\leq n\cdot\exp\left\{-\mu_{\max}\frac{N^2}{\alpha^4(n-1)^2}(3\log(3)-2)\right\} \\ &\leq n\cdot\exp\left\{-\frac{\alpha^2(n-1)}{N}\frac{N^2}{\alpha^4(n-1)^2}(3\log(3)-2)\right\} \\ &\leq n\cdot\exp\left\{-\frac{N}{\alpha^2(n-1)}\right\}.\end{aligned}$$

6 Proof of Lemma 2 (space complexity)

Proof of Lemma 2. Denote with A the event $A = \left\{\forall s \in \{1,\dots,t\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \leq \varepsilon\right\}$, we reformulate

$$\begin{aligned}&\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{t,e,j} \geq 3N \cap \left\{\forall s \in \{1,\dots,t\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \leq \varepsilon\right\}\right) \\ &= \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{t,e,j} \geq 3N \cap A\right) = \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{t,e,j} \geq 3N \,\middle|\, A\right)\mathbb{P}\left(A\right)\end{aligned}$$

While we do know that the $\widehat{z}_{t,e,j}$ are Bernoulli random variables (since they are either 0 or 1), it is not easy to compute the success probability of each $\widehat{z}_{t,e,j}$, and in addition there could be dependencies between $\widehat{z}_{t,e,j}$ and $\widehat{z}_{t,e',j'}$. Similarly to Lemma 3, we are going to find a stochastic variable to dominate $\widehat{z}_{t,e,j}$. Denoting with $u'_{s,e,j} \sim \mathcal{U}(0,1)$ a uniform random variable, we will define $w'_{s,e,j}$ as

$$w'_{s,e,j}|\mathcal{F}_{\{s,e',j'\}} = w'_{s,e,j}|\mathcal{F}_{\{s-2,m,N\}} \stackrel{\text{def}}{=\!=} \mathbb{I}\left\{u'_{s,e,j} \leq \frac{p_{t,e}}{\widetilde{p}_{s-1,e}}\right\} \sim \mathcal{B}\left(\frac{p_{t,e}}{\widetilde{p}_{s-1,e}}\right)$$

for any e' and j' such that $\{s,1,1\} \leq \{s,e',j'\} < \{s,e,j\}$. Note that $w'_{s,e,j}$, unlike $\widehat{z}_{s,e,j}$, does not have a recursive definition, and its only dependence on any other variable comes from $\widetilde{p}_{s-1,e}$. First, we peel off the last step

$$\begin{aligned}
&\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{t,e,j} \geq gN \;\middle|\; A\right) \\
&= \sum_{\mathcal{F}_{\{t-1,m,N\}}} \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\mathbb{I}\left\{u_{t,e,j} \leq \frac{\widetilde{p}_{t,e}}{\widetilde{p}_{t-1,e}}\right\}\widehat{z}_{t-1,e,j} \geq gN \;\middle|\; \mathcal{F}_{\{t-1,m,N\}} \cap A\right)\mathbb{P}\left(\mathcal{F}_{\{t-1,m,N\}} \mid A\right) \\
&\leq \sum_{\mathcal{F}_{\{t-1,m,N\}}} \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\mathbb{I}\left\{u'_{t,e,j} \leq \frac{p_{t,e}}{\widetilde{p}_{t-1,e}}\right\}\widehat{z}_{t-1,e,j} \geq gN \;\middle|\; \mathcal{F}_{\{t-1,m,N\}} \cap A\right)\mathbb{P}\left(\mathcal{F}_{\{t-1,m,N\}} \mid A\right) \\
&= \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{t,e,j}\widehat{z}_{t-1,e,j} \geq gN \;\middle|\; A\right),
\end{aligned}$$

where we used the fact that conditioned on A, Prop. 2 holds and guarantees that $\widetilde{p}_{t,e}$ is α-good, and therefore $\widetilde{p}_{t,e} \leq p_{t,e}$. Plugging this in the previous bound,

$$\begin{aligned}
\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{t,e,j} \geq gN \;\middle|\; A\right)\mathbb{P}(A) &\leq \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{t,e,j}\widehat{z}_{t-1,e,j} \geq gN \;\middle|\; A\right)\mathbb{P}(A) \\
= \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{t,e,j}\widehat{z}_{t-1,e,j} \geq gN \cap A\right) &\leq \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{t,e,j}\widehat{z}_{t-1,e,j} \geq gN\right).
\end{aligned}$$

We now proceed by peeling off layers from the end of the chain one by one. We show how to move from an iteration $s \leq t$ to $s-1$.

$$\begin{aligned}
&\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{s,e,j}\widehat{z}_{s-1,e,j} \geq gN\right) \\
&= \mathop{\mathbb{E}}_{\mathcal{F}_{\{s-2,m,N\}}}\left[\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\mathbb{I}\left\{u'_{s,e,j} \leq \frac{p_{t,e}}{\widetilde{p}_{s-1,e}}\right\}\widehat{z}_{s-1,e,j} \geq gN \;\middle|\; \mathcal{F}_{\{s-2,m,N\}}\right)\right] \\
&= \mathop{\mathbb{E}}_{\mathcal{F}_{\{s-2,m,N\}}}\left[\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\mathbb{I}\left\{u'_{s,e,j} \leq \frac{p_{t,e}}{\widetilde{p}_{s-1,e}}\right\}\mathbb{I}\left\{u_{s-1,e,j} \leq \frac{\widetilde{p}_{s-1,e}}{\widetilde{p}_{s-2,e}}\right\}\widehat{z}_{s-2,e,j} \geq gN \;\middle|\; \mathcal{F}_{\{s-2,m,N\}}\right)\right] \\
&= \mathop{\mathbb{E}}_{\mathcal{F}_{\{s-2,m,N\}}}\left[\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\mathbb{I}\left\{u'_{s-1,e,j} \leq \frac{p_{t,e}}{\widetilde{p}_{s-2,e}}\right\}\widehat{z}_{s-2,e,j} \geq gN \;\middle|\; \mathcal{F}_{\{s-2,m,N\}}\right)\right] \\
&= \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{s-1,e,j}\widehat{z}_{s-2,e,j} \geq gN\right)
\end{aligned}$$

Applying this repeatedly from $s=t$ to $s=2$ we have,

$$\begin{aligned}
&\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{t,e,j} \geq gN \cap \left\{\forall\, s \in \{1,\ldots,t\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \leq \varepsilon\right\}\right) \\
&\leq \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{t,e,j}\widehat{z}_{t-1,e,j} \geq gN\right) = \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{1,e,j}\widehat{z}_{0,e,j} \geq gN\right) = \mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}w'_{1,e,j} \geq gN\right).
\end{aligned}$$

Now, all the $w'_{1,e,j}$ are independent Bernoulli random variables, and we can bound their sum with a Hoeffding-like bound using Markov inequality,

$$\begin{aligned}
&\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t} w'_{1,e,j} \geq gN\right) = \inf_{\theta>0} \mathbb{P}\left(e^{\sum_{j=1}^{N}\sum_{e=1}^{t}\theta w'_{1,e,j}} \geq e^{\theta gN}\right) \\
&\leq \inf_{\theta>0} \frac{\mathbb{E}\left[e^{\sum_{j=1}^{N}\sum_{e=1}^{t}\theta w'_{1,e,j}}\right]}{e^{\theta gN}} = \inf_{\theta>0} \frac{\mathbb{E}\left[\prod_{j=1}^{N}\prod_{e=1}^{t} e^{\theta w'_{1,e,j}}\right]}{e^{\theta gN}} = \inf_{\theta>0} \frac{\prod_{j=1}^{N}\prod_{e=1}^{t}\mathbb{E}\left[e^{\theta w'_{1,e,j}}\right]}{e^{\theta gN}} \\
&= \inf_{\theta>0} \frac{\prod_{j=1}^{N}\prod_{e=1}^{t}(p_{t,e}e^{\theta}+(1-p_{t,e}))}{e^{\theta gN}} = \inf_{\theta>0} \frac{\prod_{j=1}^{N}\prod_{e=1}^{t}(1+p_{t,e}(e^{\theta}-1))}{e^{\theta gN}} \\
&\leq \inf_{\theta>0} \frac{\prod_{j=1}^{N}\prod_{e=1}^{t} e^{p_{t,e}(e^{\theta}-1)}}{e^{\theta gN}} \leq \inf_{\theta>0} \frac{e^{N(e^{\theta}-1)}}{e^{\theta gN}} = \inf_{\theta>0} e^{(Ne^{\theta}-N-\theta gN)},
\end{aligned}$$

where we use the fact that $1+x \leq e^x$ and by definition $w'_{1,e,j} \sim \mathcal{B}(p_{e,t})$ and $\sum_{e=1}^{t} p_{t,e} = 1$.

The choice of θ minimizing the previous expression is obtained as

$$\frac{d}{d\theta} e^{\left(Ne^{\theta}-N-\theta gN\right)} = e^{\left(Ne^{\theta}-N-\theta gN\right)}\left(Ne^{\theta}-gN\right) = 0,$$

and thus $\theta = \log(g)$. Finally,

$$\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{m,e,j} \geq gN\right) \leq \inf_{\theta} \exp\left\{N(e^{\theta}-1-\theta g)\right\} = \exp\left\{N\left(g-1-g\log(g)\right)\right\}$$

choosing $g = 3$ and plugging in the definition of N from Algorithm 1,

$$\begin{aligned}
&\mathbb{P}\left(\sum_{j=1}^{N}\sum_{e=1}^{t}\widehat{z}_{t,e,j} \geq 3N \cap \left\{\forall\, s \in \{1,\ldots,t\} : \|\widehat{\mathbf{Y}}_{\{s,m,N\}}\| \leq \varepsilon\right\}\right) \\
&\qquad \leq \exp\left\{-40\alpha^2 n \log^2(3m/\delta)/\varepsilon^2\right\} \leq \exp\left\{-\log(2m/\delta)\right\} = \frac{\delta}{2m}.
\end{aligned}$$

□

References

[1] M. B. Cohen, C. Musco, and J. Pachocki. Online row sampling. 2016. URL `http://arxiv.org/abs/1604.05448`.

[2] J. A. Kelner and A. Levin. Spectral sparsification in the semi-streaming setting. *Theory of Computing Systems*, 53(2):243–262, 2013.

[3] H. Levy. *Stochastic dominance: Investment decision making under uncertainty*. Springer, 2015.

[4] J. Pachocki. Analysis of resparsification. 2016. URL `http://arxiv.org/abs/1605.08194`.

[5] D. A. Spielman and N. Srivastava. Graph sparsification by effective resistances. *Journal on Computing*, 40(6), 2011.

[6] J. A. Tropp. Freedman's inequality for matrix martingales. *Electronic Communications in Probability*, 16:262–270, 2011.

[7] J. A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015.

[8] R. Vershynin. A note on sums of independent random matrices after Ahlswede-Winter. 2009. URL `http://www.umich.edu/~romanv/teaching/reading-group/ahlswede-winter.pdf`.